

Two-Locus Likelihoods under Variable Population Size and Fine-Scale Recombination Rate Estimation

John A. Kamm^{1,2,*}, Jeffrey P. Spence^{3,*}, Jeffrey Chan², and Yun S. Song^{1,2,4,5}

¹ Department of Statistics, University of California, Berkeley, CA 94720

² Computer Science Division, University of California, Berkeley, CA 94720

³ Computational Biology Graduate Group, University of California, Berkeley, CA 94720

⁴ Department of Integrative Biology, University of California, Berkeley, CA 94720

⁵ Department of Mathematics and Department of Biology, University of Pennsylvania, PA 19104

* These authors contributed equally to this work.

Abstract

Two-locus sampling probabilities have played a central role in devising an efficient composite likelihood method for estimating fine-scale recombination rates. Due to mathematical and computational challenges, these sampling probabilities are typically computed under the unrealistic assumption of a constant population size, and simulation studies have shown that resulting recombination rate estimates can be severely biased in certain cases of historical population size changes. To alleviate this problem, we develop here new methods to compute the sampling probability for variable population size functions that are piecewise constant. Our main theoretical result, implemented in a new software package called LDpop, is a novel formula for the sampling probability that can be evaluated by numerically exponentiating a large but sparse matrix. This formula can handle moderate sample sizes ($n \leq 50$) and demographic size histories with a large number of epochs ($\mathcal{D} \geq 64$). In addition, LDpop implements an approximate formula for the sampling probability that is reasonably accurate and scales to hundreds in sample size ($n \geq 256$). Finally, LDpop includes an importance sampler for the posterior distribution of two-locus genealogies, based on a new result for the optimal proposal distribution in the variable-size setting. Using our methods, we study how a sharp population bottleneck followed by rapid growth affects the correlation between partially linked sites. Then, through an extensive simulation study, we show that accounting for population size changes under such a demographic model leads to substantial improvements in fine-scale recombination rate estimation. LDpop is freely available for download at <https://github.com/popgenmethods/ldpop>.

1 Introduction

The coalescent with recombination (Griffiths and Marjoram, 1997) provides a basic population genetic model for recombination. For a very small number of loci and a constant population size, the likelihood (or sampling probability) can be computed via a recursion (Golding, 1984; Ethier and Griffiths, 1990; Hudson, 2001) or importance sampling (Fearnhead and Donnelly, 2001), allowing for maximum-likelihood and Bayesian estimates of recombination rates (Fearnhead and Donnelly, 2001; Hudson, 2001; McVean et al., 2002; Fearnhead et al., 2004; Fearnhead and Smith, 2005; Fearnhead, 2006).

Jenkins and Song (2009, 2010) recently introduced a new framework based on asymptotic series (in inverse powers of the recombination rate ρ) to approximate the two-locus sampling probability under a constant population size, and developed an algorithm for finding the expansion to an arbitrary order (Jenkins and Song, 2012). They also proved that only a finite number of terms in the expansion is needed to obtain the *exact* two-locus sampling probability as an analytic function of ρ . Bhaskar and Song (2012) partially extended this approach to an arbitrary number of loci and

found closed-form formulas for the first two terms in an asymptotic expansion of the multi-locus sampling distribution.

When there are more than a handful of loci, computing the sampling probability becomes intractable. A popular and tractable alternative has been to construct composite likelihoods by multiplying the two-locus likelihoods for pairs of SNPs; this pairwise composite likelihood has been used to estimate fine-scale recombination rates in humans (The International HapMap Consortium, 2007; 1000 Genomes Project Consortium, 2010), *Drosophila* (Chan et al., 2012), chimpanzees (Auton et al., 2012), microbes (Johnson and Slatkin, 2009), dogs (Auton et al., 2013), and more, and was used in the discovery of a DNA motif associated with recombination hotspots in some organisms, including humans (Myers et al., 2008), subsequently identified as a binding site of the protein PRDM9 (Myers et al., 2010; Baudat et al., 2010; Berg et al., 2010).

The pairwise composite likelihood was first suggested by Hudson (2001). The software package LDhat (McVean et al., 2004; Auton and McVean, 2007) implemented the pairwise composite likelihood and embedded it within a Bayesian MCMC algorithm for inference. Chan et al. (2012) modified this algorithm in their program LDhelmet to efficiently utilize aforementioned asymptotic formulas for the sampling probability, among other improvements. The program LDhot (Myers et al., 2005; Auton et al., 2014) uses the composite likelihood as a test statistic to detect recombination hotspots, in conjunction with coalescent simulation to determine appropriate null distributions.

Because of mathematical and computational challenges, LDhat, LDhelmet, and LDhot all assume a constant population size model to compute the two-locus sampling probabilities. This is an unrealistic assumption, and it would be desirable to account for known demographic events, such as bottlenecks or population growth. Previous studies (McVean et al., 2002; Chan et al., 2012; Smith and Fearnhead, 2005) have shown that incorrectly assuming constant population size can lead these composite-likelihood methods to produce biased estimates. Furthermore, Johnston and Cutler (2012) observed that a sharp bottleneck followed by rapid growth can lead LDhat to infer spurious recombination hotspots if it assumes a constant population size.

Hudson (2001) proposed Monte Carlo computation of two locus likelihoods by simulating genealogies. While this generalizes to arbitrarily complex demographics, it would be desirable to have a deterministic formula, as naive Monte Carlo computation sometimes has difficulty approximating small probabilities.

In this paper, we show how to compute the two-locus sampling probability exactly under variable population size histories that are piecewise constant. Our approach relies on the Moran model (Moran, 1958; Ethier and Kurtz, 1993; Donnelly and Kurtz, 1999), a stochastic process closely related to the coalescent. We have implemented our results in a freely available software package, LDpop, that efficiently produces lookup tables of two-locus likelihoods under variable population size. These lookup tables can then be used by other programs that use composite likelihoods to infer recombination maps.

Our main result is an exact formula, introduced in Theorem 1, that involves exponentiating sparse m -by- m matrices containing $O(m)$ nonzero entries, where $m = O(n^6)$ under a bi-allelic model, with n being the sample size. We derive this formula by constructing a Moran-like process in which sample paths can be coupled with the two-locus coalescent, and by applying a reversibility argument.

Theorem 1 has a high computational cost, and our implementation in LDpop can practically handle low to moderate sample sizes ($n < 50$) on a 24-core compute server. We have thus implemented an approximate formula that is much faster, and scales to sample sizes in the hundreds. This formula is computed by exponentiating a sparse matrix with $O(n^3)$ nonzero entries, and is based on a previous two-locus Moran process (Ethier and Kurtz, 1993), which we have implemented

and extended to the case of variable population size. While this formula does not give the exact likelihood, it provides a reasonable approximation, and converges to the true value in an appropriate limit.

In addition to these exact and approximate formulas, LDpop also includes a highly efficient importance sampler for the posterior distribution of two-locus genealogies. This can be used to infer the genealogy at a pair of sites, and also provides an alternative method for computing two-locus likelihoods. Our importance sampler is based on an optimal proposal distribution that we characterize in Theorem 2. It generalizes previous results for the constant size case, which have been used to construct importance samplers for both the single-population, two-locus case (Fearnhead and Donnelly, 2001; Dialdestoro et al., 2016) and for other constant-demography coalescent scenarios (Stephens and Donnelly, 2000; De Iorio and Griffiths, 2004; Griffiths et al., 2008; Hobolth et al., 2008; Jenkins, 2012; Koskela et al., 2015). The key ideas of Theorem 2 should similarly generalize to other contexts of importance sampling a time-inhomogeneous coalescent.

Using a simulation study, we show that using LDpop to account for demography substantially improves the composite-likelihood inference of recombination maps. We also use LDpop to gain a qualitative understanding of linkage disequilibrium by examining the r^2 statistic. Finally, we examine how LDpop scales in terms of sample size n and the number \mathcal{D} of demographic epochs. The exact formula can handle n in the tens, while the approximate formula can handle n in the hundreds. Additionally, we find that the runtime of LDpop is not very sensitive to \mathcal{D} , so LDpop can handle size histories with a large number of pieces.

Software availability: Our software package LDpop is freely available for download at <https://github.com/popgenmethods/ldpop>.

2 Background

Here we describe our notational convention, and review some key concepts regarding the coalescent with recombination and the two-locus Moran model.

2.1 Notation

Let $\frac{\theta}{2}$ denote the mutation rate per locus per unit time, $\mathbf{P} = (P_{ij})_{i,j \in \mathcal{A}}$ the transition probabilities between alleles given a mutation, and $\mathcal{A} = \{0, 1\}$ the set of alleles (our formulas can be generalized to $|\mathcal{A}| > 2$, but this increases the computational complexity). Let $\frac{\rho}{2}$ denote the recombination rate per unit time. We consider a single panmictic population, with piecewise-constant effective population sizes. In particular, we assume D pieces, with endpoints $-\infty = t_{-D} < t_{-D+1} < \dots < t_{-1} < t_0 = 0$, where 0 corresponds to the present and $t < 0$ corresponds to a time in the past. The piece $(t_d, t_{d+1}]$ is assumed to have scaled population size η_d . Going backwards in time, two lineages coalesce (find a common ancestor) at rate $\frac{1}{\eta_d}$ within the interval $(t_d, t_{d+1}]$.

We allow the haplotypes to have missing (unobserved) alleles at each locus, and use $*$ to denote such alleles. We denote each haplotype as having type a , b , or c , where a haplotypes are only observed at the left locus, b haplotypes are only observed at the right locus, and c haplotypes are observed at both loci. Overloading notation, we sometimes refer to the left locus as the a locus, and the right locus as the b locus. We use $\mathbf{n} = \{n_{i*}, n_{*j}, n_{k\ell}\}_{i,j,k,\ell \in \mathcal{A}}$ to denote the configuration of an unordered collection of two-locus haplotypes, with n_{i*} denoting the number of a -types with allele i , n_{*j} denoting the number of b -types with allele j , and $n_{k\ell}$ denoting the number of c -types with alleles k and ℓ .

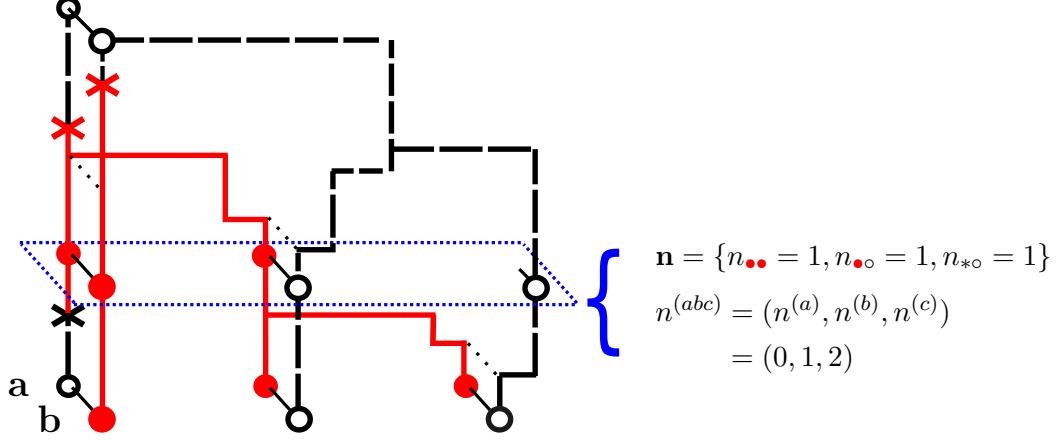


Figure 1: An ancestral recombination graph (ARG) at two loci, labeled a and b , each with two alleles (\bullet and \circ). The notation $\mathbf{n}, n_{ij}, n^{(abc)}$ are illustrated for the configuration between the first coalescence and second recombination event.

Table 1: Backward in time transition rates of $n_t^{(abc)} = (n_t^{(a)}, n_t^{(b)}, n_t^{(c)})$ within time interval $(t_d, t_{d+1}]$ under the coalescent with recombination.

End state	Rate
$n_t^{(a)} + 1, n_t^{(b)} + 1, n_t^{(c)} - 1$	$\frac{\rho}{2} n_t^{(c)}$
$n_t^{(a)} - 1, n_t^{(b)}, n_t^{(c)}$	$\frac{1}{\eta_d} n_t^{(a)} \left(\frac{n_t^{(a)} - 1}{2} + n_t^{(c)} \right)$
$n_t^{(a)}, n_t^{(b)} - 1, n_t^{(c)}$	$\frac{1}{\eta_d} n_t^{(b)} \left(\frac{n_t^{(b)} - 1}{2} + n_t^{(c)} \right)$
$n_t^{(a)}, n_t^{(b)}, n_t^{(c)} - 1$	$\frac{1}{\eta_d} \binom{n_t^{(c)}}{2}$
$n_t^{(a)} - 1, n_t^{(b)} - 1, n_t^{(c)} + 1$	$\frac{1}{\eta_d} n_t^{(a)} n_t^{(b)}$

Suppose \mathbf{n} has $n^{(abc)} = (n^{(a)}, n^{(b)}, n^{(c)})$ haplotypes of type a, b, c respectively. We define the *sampling probability* $\mathbb{P}_t(\mathbf{n})$ to be the probability of sampling \mathbf{n} at time t , given that we observed $n^{(a)}, n^{(b)}, n^{(c)}$ haplotypes of type a, b, c , under the coalescent with recombination (next subsection).

2.2 The ARG and the coalescent with recombination

The Ancestral Recombination Graph (ARG) is the multi-locus genealogy relating a sample (Figure 1). The coalescent with recombination (Griffiths, 1991) gives the limiting distribution of the ARG under a wide class of population models, including the Wright-Fisher model and the Moran model.

Let $n_t^{(c)}$ be the number of lineages at time t that are ancestral to the observed present-day sample at both loci. Similarly, let $n_t^{(a)}$ and $n_t^{(b)}$ be the number of lineages that are ancestral at only the a or b locus, respectively. Under the coalescent with recombination, $n_t^{(abc)} = (n_t^{(a)}, n_t^{(b)}, n_t^{(c)})$ is a backward in time Markov chain, where each c type lineage splits (recombines) into one a and one b lineage at rate $\frac{\rho}{2}$, and each pair of lineages coalesces at rate $\frac{1}{\eta_d}$ within the time interval $(t_d, t_{d+1}]$.

Table 1 gives the transition rates of $n_t^{(abc)}$.

After sampling the history of coalescence and recombination events $\{n_t^{(abc)}\}_{t \leq 0}$, we drop mutations down at rate $\frac{\theta}{2}$ per locus, with alleles mutating according to \mathbf{P} , and the alleles of the common




Figure 2: A finite two-locus Moran model with $N = 4$ particles. Each lineage copies itself onto every other lineage at rate $\frac{1}{2\eta}$. Mutations arise at rate $\frac{\theta}{2}$ per allele per locus. Recombination follows dynamics from Ethier and Kurtz (1993): every pair of lineages experience a crossover recombination at rate $\frac{\rho}{2(N-1)}$. Here, the 2nd and 3rd lineages swap their b alleles through a crossover. The sampling probability for this model agrees with the coalescent at each locus marginally, but not jointly at both loci (though the discrepancy disappears as $N \rightarrow \infty$).

ancestor assumed to be at the stationary distribution. This gives us a sample path $\{\mathbf{n}_t\}_{t \leq 0}$, where \mathbf{n}_0 is the observed sample at the present, and \mathbf{n}_t is the collection of ancestral haplotypes at time t . Under this notation, the sampling probability at time t is defined as

$$\mathbb{P}_t(\mathbf{n}) := \mathbb{P}(\mathbf{n}_t = \mathbf{n} \mid n_t^{(abc)} = n^{(abc)}). \quad (1)$$

2.3 Two-locus Moran model

The Moran model is a stochastic process closely related to the coalescent, and plays a central role in our results. Here, we review a two-locus Moran model with recombination dynamics from Ethier and Kurtz (1993). We note that there are multiple versions of the two-locus Moran model, and in particular Donnelly and Kurtz (1999) describe a Moran model with different recombination dynamics.

The Moran model with N lineages is a finite population model evolving forward in time. In particular, let \mathbf{M}_t denote a collection of N two-locus haplotypes at time t (with no missing alleles). Then \mathbf{M}_t is a Markov chain going forwards in time that changes due to mutation, recombination, and copying events.

Let $\mathbf{\Lambda}_{(N)}^d$ denote the transition matrix of \mathbf{M}_t within $(t_d, t_{d+1}]$. We describe the rates of $\mathbf{\Lambda}_{(N)}^d$. For the mutation events, each allele mutates at rate $\frac{\theta}{2}$ according to transition matrix \mathbf{P} . For the copying events, each lineage of \mathbf{M}_t copies its haplotype onto each other lineage at rate $\frac{1}{2\eta_d}$ within the time interval $(t_d, t_{d+1}]$. Biologically, this corresponds to one lineage dying out, and being replaced by the offspring of another lineage, which occurs more frequently when genetic drift is high (i.e. when the population size η_d is small). Finally, every pair of lineages in \mathbf{M}_t swap genetic material through a crossover recombination at rate $\frac{\rho}{2(N-1)}$. A crossover between haplotypes (i_1, j_1) and (i_2, j_2) results in new haplotypes (i_1, j_2) and (i_2, j_1) , and the configuration resulting from the crossover is $\mathbf{M}_t - \mathbf{e}_{i_1, j_1} - \mathbf{e}_{i_2, j_2} + \mathbf{e}_{i_1, j_2} + \mathbf{e}_{i_2, j_1}$. See Figure 2 for illustration.

Let $\mathbb{P}_t^{(N)}(\mathbf{n})$ be the probability of sampling \mathbf{n} at time t under this Moran model, for a configura-

tion \mathbf{n} with sample size $n \leq N$. $\mathbb{P}_t^{(N)}(\mathbf{n})$ is given by first sampling $\mathbf{M}_{\min(t, t_{-D+1})}$ from the stationary distribution $\boldsymbol{\lambda}_{(N)}^{-D}$ of $\boldsymbol{\Lambda}_{(N)}^{-D}$, then propagating $\{\mathbf{M}_s\}_{s \leq 0}$ forward-in-time to t , and then sampling \mathbf{n} without replacement from \mathbf{M}_t . So,

$$\begin{aligned} [\mathbb{P}_{(N)}(\mathbf{M}_t = \mathbf{M})]_{\mathbf{M}} &= \boldsymbol{\lambda}_{(N)}^{-D} \prod_{d=-D+1}^{-1} e^{\boldsymbol{\Lambda}_{(N)}^d [\min(t, t_{d+1}) - \min(t, t_d)]} \\ \mathbb{P}_t^{(N)}(\mathbf{n}) &= \sum_{\mathbf{M}} \mathbb{P}_{(N)}(\mathbf{M}_t = \mathbf{M}) \mathbb{P}(\mathbf{n} \mid \mathbf{M}), \end{aligned} \quad (2)$$

where $\mathbb{P}(\mathbf{n} \mid \mathbf{M})$ denotes the probability of sampling \mathbf{n} without replacement from \mathbf{M} , and $[\mathbb{P}_{(N)}(\mathbf{M}_t = \mathbf{M})]_{\mathbf{M}}$ and $\boldsymbol{\lambda}_{(N)}^{-D}$ are row vectors here.

In general, $\mathbb{P}_t^{(N)}(\mathbf{n}) \neq \mathbb{P}_t(\mathbf{n})$, so \mathbf{M}_t disagrees with the coalescent with recombination. However, the likelihood under \mathbf{M}_t converges to the correct value, $\mathbb{P}_t^{(N)}(\mathbf{n}) \rightarrow \mathbb{P}_t(\mathbf{n})$ as $N \rightarrow \infty$. In fact, even for $N = n = 20$, we find that $\mathbb{P}_t^{(N)}(\mathbf{n})$ provides a reasonable approximation for practical purposes (Sections 4.2, 5.2). We refer to (2), i.e. the likelihood under \mathbf{M}_t , as the “approximate likelihood formula”, in contrast to the exact formula we present in Theorem 1 below. This approximate formula is included in LDpop as a faster, more scalable alternative to the exact formula of Theorem 1.

3 Theoretical Results

In this section, we describe our theoretical results. Proofs are deferred to Appendix D.

3.1 Exact formula for the sampling probability

Our main result is an explicit formula for the sampling probability $\mathbb{P}(\mathbf{n})$, presented in Theorem 1. We present an outline here; the proof is delayed to Appendix D.1.

The idea is to construct a forward-in-time Markov process $\tilde{\mathbf{M}}_t$, and relate its distribution to the coalescent with recombination. $\tilde{\mathbf{M}}_t$ is similar to the Moran model \mathbf{M}_t in Section 2.3, except that $\tilde{\mathbf{M}}_t$ allows partially specified a - and b -types, whereas all lineages in \mathbf{M}_t are fully specified c -types. Specifically, the state space of $\tilde{\mathbf{M}}_t$ is $\mathcal{N} = \{\mathbf{n} : n^{(abc)} = (k, k, n - k), 0 \leq k \leq n\}$ the collection of sample configurations with n specified (non-missing) alleles at each locus. The state of $\tilde{\mathbf{M}}_t$ changes due to copying, mutation, recombination, and “recoalescence” events. Copying and mutation dynamics are similar to \mathbf{M}_t , but recombination is different: every c -type splits into a - and b -types at rate $\frac{\rho}{2}$, and every pair of a - and b -types “recoalesces” back into a c -type at rate $\frac{1}{\eta_d}$. An illustration is shown in Figure 3b; $\tilde{\mathbf{M}}_t$ is described in more detail in Appendix D.1. Within interval $(t_d, t_{d+1}]$, we denote the transition rate matrix of $\tilde{\mathbf{M}}_t$ as $\tilde{\boldsymbol{\Lambda}}^d$, a square matrix indexed by \mathcal{N} , with entries given in Table 2.

$\tilde{\mathbf{M}}_t$ does not itself yield the correct sampling probability. The basic issue is that $\tilde{\mathbf{M}}_t$ has c -types splitting into a - and b -types at rate $\frac{\rho}{2}$ going *forwards in time*, but under the coalescent with recombination, this needs to happen at rate $\frac{\rho}{2}$ going *backwards in time*. Similarly, pairs of a - and b -types merge into a single c -type at rate $\frac{1}{\eta_d}$ going forwards in time under $\tilde{\mathbf{M}}_t$, but going backwards in time under the coalescent with recombination.

However, it is possible to “reverse” the direction of the recombinations ($c \rightarrow a, b$) and recoalescences ($a, b \rightarrow c$), to get a new process that does match the two-locus coalescent. In particular, let C_t be the number of c types in $\tilde{\mathbf{M}}_t$, illustrated in Figure 3a. Then C_t is a reversible Markov chain, whose rate matrix in $(t_d, t_{d+1}]$ is $\boldsymbol{\Gamma}^d$ a tridiagonal square matrix indexed by $\{0, 1, \dots, n\}$, with




Figure 3: The process $\{\tilde{\mathbf{M}}_t\}_{t \leq 0}$ with rates $\tilde{\Lambda}^d$ (Table 2) used to prove Theorem 1. This process is similar to $\{\mathbf{M}_t\}$ (Figure 2), in that n is fixed, and alleles change due to copying and mutation. However $\tilde{\mathbf{M}}_t$ allows partially specified a - and b -types, with c -types recombining into a pair of a - and b -types, and pairs of a - and b -types “recoalescing” into c -types. (a) The process $\{C_t\}_{t \leq 0}$ of just recombination ($c \rightarrow a, b$) and “recoalescence” ($a, b \rightarrow c$) events. (b) The full process $\{\tilde{\mathbf{M}}_t\}_{t \leq 0}$ including copying, mutation, recombination, and “recoalescence” events.

Table 2: Nonzero entries of the rate matrix $\tilde{\Lambda}^d$ for the interval $(t_d, t_{d+1}]$.

\mathbf{m}	$\tilde{\Lambda}_{\mathbf{n}, \mathbf{m}}^d$
$\mathbf{n} - \mathbf{e}_{i*} + \mathbf{e}_{j*}$	$\frac{1}{\eta_d} n_{i*} (\frac{1}{2} n_{j*} + \sum_{k \in \mathcal{A}} n_{jk}) + \frac{\theta}{2} P_{ij} n_{i*}$
$\mathbf{n} - \mathbf{e}_{*i} + \mathbf{e}_{*j}$	$\frac{1}{\eta_d} n_{*i} (\frac{1}{2} n_{*j} + \sum_{k \in \mathcal{A}} n_{kj}) + \frac{\theta}{2} P_{ij} n_{*i}$
$\mathbf{n} - \mathbf{e}_{ij} + \mathbf{e}_{kl}$	$\frac{1}{2\eta_d} n_{ij} n_{kl} + \frac{\theta}{2} (\delta_{ik} P_{jl} + \delta_{jl} P_{ik}) n_{ij}$
$\mathbf{n} - \mathbf{e}_{ij} + \mathbf{e}_{i*} + \mathbf{e}_{*j}$	$\frac{\rho}{2} n_{ij}$
$\mathbf{n} - \mathbf{e}_{i*} - \mathbf{e}_{*j} + \mathbf{e}_{ij}$	$n_{i*} n_{*j} \frac{1}{\eta_d}$
\mathbf{n}	$-\frac{1}{\eta_d} (n^{(a)} + n^{(b)} + n^{(c)}) - \frac{\rho}{2} n^{(c)} - \frac{\theta}{2} \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{A} \cup \{*\}} (n_{ij} + n_{ji})$

$\Gamma_{m, m-1}^d = \frac{\rho}{2} m$, $\Gamma_{m, m+1}^d = (n - m)^2 \frac{1}{\eta_d}$, and $\Gamma_{m, m}^d = -\Gamma_{m, m-1}^d - \Gamma_{m, m+1}^d$. Exploiting the reversibility of C_t allows us to relate the distribution of $\tilde{\mathbf{M}}_t$ to the coalescent with recombination, and obtain the following result.

Theorem 1. Let $(\gamma_0^d, \dots, \gamma_n^d)$ be the stationary distribution of Γ^d , and let the row vector $\tilde{\gamma}^d$ be indexed by \mathcal{N} , with $\tilde{\gamma}_{\mathbf{n}}^d = \gamma_m^d$ if \mathbf{n} has m lineages of type c . Denote the stationary distribution of $\tilde{\Lambda}^d$ by the row vector $\tilde{\lambda}^d = (\lambda_{\mathbf{n}}^d)_{\mathbf{n} \in \mathcal{N}}$. Let \odot and \div denote component-wise multiplication and division, and recursively define the row vector $\mathbf{p}^d = (p_{\mathbf{n}}^d)_{\mathbf{n} \in \mathcal{N}}$ by

$$\begin{aligned} \mathbf{p}^{-D+1} &= \tilde{\lambda}^{-D} \div \tilde{\gamma}^{-D} \\ \mathbf{p}^{d+1} &= [(\mathbf{p}^d \odot \tilde{\gamma}^d) e^{\tilde{\Lambda}^d(t_{d+1}-t_d)}] \div \tilde{\gamma}^d. \end{aligned} \quad (3)$$

Then, for $\mathbf{n} \in \mathcal{N}$, we have $\mathbb{P}_0(\mathbf{n}) = p_{\mathbf{n}}^0$.

Note that Theorem 1 gives $\mathbb{P}_0(\mathbf{n})$ for $\mathbf{n} \in \mathcal{N}$. This includes all fully specified \mathbf{n} , i.e. with $n^{(abc)} = (0, 0, n)$, and suffices for the application considered in Section 4.2. If necessary, $\mathbb{P}_0(\mathbf{n})$ for

partially specified \mathbf{n} can be computed by summing over the fully specified configurations consistent with \mathbf{n} .

For $|\mathcal{A}| = 2$ alleles per locus, $\tilde{\mathbf{\Lambda}}^d$ is an $O(n^6) \times O(n^6)$ matrix, so naively computing the matrix multiplication in Theorem 1 would cost $O(n^{12})$ time. However, $\tilde{\mathbf{\Lambda}}^d$ is sparse, with $O(n^6)$ nonzero entries, allowing efficient algorithms to compute Theorem 1 (up to numerical precision) in $O(n^6 \mathcal{T})$ time, where \mathcal{T} is some finite number of matrix-vector multiplications. See Appendix C.1 and C.2 for more details.

3.2 Importance sampling

In addition to the approximate (2) and exact (3) formulas for the sampling probability, LDpop includes an importance sampler for the two-locus ARG $\mathbf{n}_{\leq 0} = \{\mathbf{n}_t\}_{t \leq 0}$. This provides a method to sample from the posterior distribution of two-locus ARGs, and also provides an alternative method for computing $\mathbb{P}_0(\mathbf{n})$. This importance sampler is based on Theorem 2 below, which characterizes the optimal proposal distribution for the two-locus coalescent with recombination under variable population size.

Let the proposal distribution $Q(\mathbf{n}_{\leq 0})$ be a probability distribution on $\{\mathbf{n}_{\leq 0} : \mathbf{n}_0 = \mathbf{n}\}$ whose support contains that of $\mathbb{P}(\mathbf{n}_{\leq 0} \mid \mathbf{n}_0 = \mathbf{n})$. Then we have

$$\mathbb{P}_0(\mathbf{n}) = \int_{\mathbf{n}_{\leq 0} : \mathbf{n}_0 = \mathbf{n}} \frac{d\mathbb{P}(\mathbf{n}_{\leq 0})}{dQ(\mathbf{n}_{\leq 0})} dQ(\mathbf{n}_{\leq 0}),$$

and so, if $\mathbf{n}_{\leq 0}^{(1)}, \dots, \mathbf{n}_{\leq 0}^{(K)} \sim Q$ i.i.d., the sum

$$\frac{1}{K} \sum_{k=1}^K \frac{d\mathbb{P}(\mathbf{n}_{\leq 0}^{(k)})}{dQ(\mathbf{n}_{\leq 0}^{(k)})} \tag{4}$$

converges almost surely to $\mathbb{P}_0(\mathbf{n})$ as $K \rightarrow \infty$ by the Law of Large Numbers. Hence, (4) provides a Monte Carlo approximation to $\mathbb{P}_0(\mathbf{n})$. The optimal proposal is the posterior distribution $Q_{\text{opt}}(\mathbf{n}_{\leq 0}) = \mathbb{P}(\mathbf{n}_{\leq 0} \mid \mathbf{n}_0)$, for then (4) is exactly

$$\frac{1}{K} \sum_{k=1}^K \frac{d\mathbb{P}(\mathbf{n}_{\leq 0}^{(k)})}{d\mathbb{P}(\mathbf{n}_{\leq 0} \mid \mathbf{n}_0)} = \frac{1}{K} \sum_{k=1}^K \frac{d\mathbb{P}(\mathbf{n}_{\leq 0}^{(k)})}{d\mathbb{P}(\mathbf{n}_{\leq 0}) / \mathbb{P}(\mathbf{n}_0)} = \mathbb{P}(\mathbf{n}_0),$$

even for $K = 1$.

The following theorem, which we prove in Appendix D.2, characterizes the optimal posterior distribution $Q_{\text{opt}}(\mathbf{n}_{\leq 0}) = \mathbb{P}(\mathbf{n}_{\leq 0} \mid \mathbf{n}_0)$ for variable population size:

Theorem 2. *The process $\{\mathbf{n}_t\}_{t \leq 0}$ is a backward-in-time Markov chain with inhomogeneous rates, whose rate matrix at time t is given by*

$$q_{\mathbf{n}, \mathbf{m}}^{(t)} = \begin{cases} \phi_{\mathbf{n}, \mathbf{m}}^{(t)} \frac{\mathbb{P}_t(\mathbf{m})}{\mathbb{P}_t(\mathbf{n})}, & \text{if } \mathbf{m} \neq \mathbf{n}, \\ \phi_{\mathbf{n}, \mathbf{n}}^{(t)} - \frac{d}{dt} \log \mathbb{P}_t(\mathbf{n}), & \text{if } \mathbf{m} = \mathbf{n}, \end{cases}$$

where $\phi^{(t)} = (\phi_{\mathbf{n}, \mathbf{m}}^{(t)})$ is a square matrix, indexed by configurations \mathbf{n} , with entries given by Table 3

Table 3: Nonzero entries of the $\phi^{(t)}$ matrix of Theorem 2, for $t \in (t_d, t_{d+1}]$.

\mathbf{m}	$\phi_{\mathbf{n}, \mathbf{m}}^{(t)}$
$\mathbf{n} - \mathbf{e}_{i*} + \mathbf{e}_{j*}$	$\frac{\theta}{2} P_{ji}(n_{j*} + 1)$
$\mathbf{n} - \mathbf{e}_{*i} + \mathbf{e}_{*j}$	$\frac{\theta}{2} P_{ji}(n_{*j} + 1)$
$\mathbf{n} - \mathbf{e}_{ij} + \mathbf{e}_{kl}$	$\frac{\theta}{2} (\delta_{ik} P_{lj} + \delta_{jl} P_{ki})(n_{kl} + 1)$
$\mathbf{n} - \mathbf{e}_{ij} + \mathbf{e}_{i*} + \mathbf{e}_{*j}$	$\frac{\rho}{2} n^{(c)}(n_{i*} + 1)(n_{*j} + 1)$
$\mathbf{n} - \mathbf{e}_{ij}$	$\frac{1}{\eta_d} \binom{n^{(c)}}{2} (n_{ij} - 1)$
$\mathbf{n} - \mathbf{e}_{i*}$	$\frac{1}{\eta_d} \left[\binom{n^{(a)}}{2} (n_{i*} - 1) + n^{(a)} n^{(c)} \sum_j n_{ij} \right]$
$\mathbf{n} - \mathbf{e}_{*i}$	$\frac{1}{\eta_d} \left[\binom{n^{(b)}}{2} (n_{*i} - 1) + n^{(b)} n^{(c)} \sum_j n_{ji} \right]$
\mathbf{n}	$-\frac{1}{\eta_d} \binom{n}{2} - \frac{\rho}{2} n^{(c)} - \frac{\theta}{2} \sum_i (1 - P_{ii}) [n_{i*} + n_{*i} + \sum_j (n_{ij} + n_{ji})]$

and equal to

$$\begin{aligned} \phi_{\mathbf{n}, \mathbf{m}}^{(t)} = & \frac{d}{ds} \left[\mathbb{P}(n_{t-s}^{(abc)} = m^{(abc)} \mid n_t^{(abc)} = n^{(abc)}) \right. \\ & \left. \times \mathbb{P}(\mathbf{n}_t = \mathbf{n} \mid \mathbf{n}_{t-s} = \mathbf{m}, n_t^{(abc)} = n^{(abc)}) \right] \Big|_{s=0}. \end{aligned} \quad (5)$$

Intuitively, the matrix $\phi^{(t)}$ in Table 3 is a linear combination of two rate matrices, one for propagating $n_t^{(abc)}$ an infinitesimal distance backwards in time, and another for propagating \mathbf{n}_t an infinitesimal distance forward in time. This is because $n_t^{(abc)}$ is generated by sampling coalescent and recombination events backwards in time, and then \mathbf{n}_t is generated by dropping mutations on the ARG and propagating the allele values forward in time.

Theorem 2 generalizes previous results for the optimal proposal distribution in the constant size case (Stephens and Donnelly, 2000; Fearnhead and Donnelly, 2001). In that case, the conditional probability of the parent \mathbf{m} of \mathbf{n} is $\phi_{\mathbf{n}, \mathbf{m}} \frac{\mathbb{P}(\mathbf{m})}{\mathbb{P}(\mathbf{n})}$. Note the constant size case is time-homogeneous, so the dependence on t is dropped, and the waiting times between events in the ARG are not sampled (i.e., only the embedded jump chain of $\mathbf{n}_{\leq 0}$ is sampled).

We construct our proposal distribution $\hat{Q}(\mathbf{n}_{\leq 0})$ by approximating the optimal proposal distribution $Q_{\text{opt}}(\mathbf{n}_{\leq 0}) = \mathbb{P}(\mathbf{n}_{\leq 0} \mid \mathbf{n}_0)$. This requires approximating the rate $q_{\mathbf{n}, \mathbf{m}}^{(t)} = \phi_{\mathbf{n}, \mathbf{m}}^{(t)} \frac{\mathbb{P}_t(\mathbf{m})}{\mathbb{P}_t(\mathbf{n})}$. We use the approximation $\hat{q}_{\mathbf{n}, \mathbf{m}}^{(t)} = \phi_{\mathbf{n}, \mathbf{m}}^{(t)} \frac{\mathbb{P}_t^{(N)}(\mathbf{m})}{\mathbb{P}_t^{(N)}(\mathbf{n})}$, with $\mathbb{P}_t^{(N)}(\mathbf{n})$ from the approximate likelihood formula (2). To save computation, we only compute $\mathbb{P}_t^{(N)}(\mathbf{n})$ along a grid of time points, and then linearly interpolate $\hat{q}_{\mathbf{n}, \mathbf{m}}^{(t)}$ between the points. See Appendix A for more details on our proposal distribution \hat{Q} .

As detailed in Section 5.3, \hat{Q} is a highly efficient proposal distribution, typically yielding effective sample sizes (ESS) between 80 to 100% per sample for the demography and ρ values we considered.

4 Application

Previous simulation studies (McVean et al., 2002; Chan et al., 2012; Smith and Fearnhead, 2005) have shown that if the demographic model is misspecified, composite-likelihood methods (which so far have assumed a constant population size) can produce recombination rate estimates that are biased. Many populations, including that of humans and *D. melanogaster*, have undergone bottlenecks and expansions in the recent past (Gutenkunst et al., 2009; Choudhary and Singh,

1987), and it has been argued (Johnston and Cutler, 2012) that such demographies can severely affect recombination rate estimation, and can cause the appearance of spurious recombination hotspots.

In this section, we apply our software LDpop to show that accounting for demography improves fine-scale recombination rate estimation. We first examine how a population bottleneck followed by rapid growth affects the correlation between partially linked sites. We then study composite likelihood estimation of recombination maps under a population bottleneck. We find that accounting for demography with either the exact (Theorem 1) or approximate (2) likelihood formula substantially improves accuracy. Furthermore, this improvement is robust to minor misspecification of the demography due to not knowing the true demography in practice.

Throughout this section, we use an example demography with $D = 3$ epochs, consisting of a sharp population bottleneck followed by a rapid expansion. Specifically, the population size history $\eta(t)$, in coalescent-scaled units, is given by

$$\eta(t) = \begin{cases} 100, & -0.5 < t \leq 0, \\ 0.1, & -0.58 < t \leq -0.5, \\ 1, & t \leq -0.58. \end{cases} \quad (6)$$

Under this model and $n = 2$, the expected time of common ancestor is $\mathbb{E}[T_{\text{MRCA}}] \approx 1$. We thus compare this demography against a constant size demography with coalescent-scaled size of $\eta \equiv 1$, as this is the population size that would be estimated using the pairwise heterozygosity (Tajima, 1983). We use a coalescent-scaled mutation rate of $\frac{\theta}{2} = 0.004$ per base, which is roughly the mutation rate of *D. melanogaster* (Chan et al., 2012).

While the size history $\eta(t)$ of (6) is fairly simple, with only $\mathcal{D} = 3$ epochs, we stress that LDpop can in fact handle much more complex size histories. For example, in Section 5, we show that LDpop can easily handle a demography with $\mathcal{D} = 64$, with little additional cost in runtime.

4.1 Linkage disequilibrium and two-locus likelihoods

One statistic of linkage disequilibrium is

$$\hat{r}^2 = \frac{[\hat{x}_{11} - \hat{x}_1^{(a)}\hat{x}_1^{(b)}]^2}{\hat{x}_0^{(a)}\hat{x}_1^{(a)}\hat{x}_0^{(b)}\hat{x}_1^{(b)}},$$

where $\hat{x}_{ij} = \frac{n_{ij}}{n}$ is the fraction of the sample with haplotype ij , with $\hat{x}_i^{(a)} = \sum_j \hat{x}_{ij}$ and $\hat{x}_j^{(b)} = \sum_i \hat{x}_{ij}$. In words, \hat{r}^2 is the sample square-correlation of a random allele at locus a with a random allele at locus b . We let $r^2 = \lim_{n \rightarrow \infty} \hat{r}^2$ denote the population square-correlation. There has been considerable theoretical interest in understanding moments of r^2 , \hat{r}^2 , and related statistics (Ohta and Kimura, 1969; Maruyama, 1982; Hudson, 1985; McVean, 2002; Song and Song, 2007). Additionally, $n\hat{r}^2$ approximately follows a χ_1^2 distribution when $r^2 = 0$, which provides a test for the statistical significance of linkage disequilibrium (Weir, 1996, p. 113).

Using LDpop, we can compute the distribution of \hat{r}^2 for piecewise constant models. In Figure 4, we show $\mathbb{E}[\hat{r}^2]$ for a sample size $n = 20$ under the 3-epoch model (6) and the constant population size model. Under the 3-epoch demography, $\mathbb{E}[\hat{r}^2]$ is much lower for small ρ and decays more rapidly as $\rho \rightarrow \infty$. In other words, the constant model requires higher ρ to break down LD to the same level, which suggests that incorrectly assuming a constant demography will lead to upward-biased estimates of the recombination rate (as pointed out by an anonymous reviewer). We confirm this




Figure 4: Expected linkage disequilibrium $\mathbb{E}[\hat{r}^2]$ for the 3-epoch model in (6) and the constant size model with $\eta \equiv 1$, as a function of recombination rate ρ for a sample size of $n = 20$. Under the 3-epoch model, even nearby sites are expected to be quite uncorrelated.

in Section 4.2.

4.2 Fine-scale recombination rate estimation

For a sample of n haplotypes observed at L SNPs, let $\mathbf{n}[a, b]$ be the two-locus sample observed at SNPs $a, b \in \{1, \dots, L\}$, and let $\rho[a, b]$ be the recombination rate between SNPs a and b . The programs LDhat (McVean et al., 2002, 2004; Auton and McVean, 2007), LDhot (Myers et al., 2005; Auton et al., 2014), and LDhelmet (Chan et al., 2012) infer hotspots and recombination maps ρ using the following composite likelihood due to Hudson (2001):

$$\prod_{a, b: 0 < b - a < W} \mathbb{P}(\mathbf{n}[a, b]; \rho[a, b]) \quad (7)$$

where W denotes some window size in which to consider pairs of sites (a finite window size W removes the computational burden and statistical noise from distant, uninformative sites (Fearnhead, 2003; Smith and Fearnhead, 2005)).

We used LDpop to generate four likelihood tables, which we then used with LDhat and LDhelmet to estimate recombination maps for simulated data. The four tables, denoted by $\mathcal{L}_{\text{const}}$, $\mathcal{L}_{\text{exact}}$, $\mathcal{L}_{\text{approx}}$, and $\mathcal{L}_{\text{miss}}$, are defined as follows:

1. (“Constant”) $\mathcal{L}_{\text{const}}$ denotes a likelihood table that assumes a constant population size of $\eta \equiv 1$.
2. (“Exact”) $\mathcal{L}_{\text{exact}}$ denotes a likelihood table that assumes the correct 3-epoch population size history η defined in (6).
3. (“Approximate”) $\mathcal{L}_{\text{approx}}$ denotes a likelihood table with the correct size history η in (6), but using the (much faster) approximate likelihood formula (2) with $N = n$ Moran particles.

4. (“Misspecified”) $\mathcal{L}_{\text{miss}}$ denotes a likelihood table that assumes a misspecified demography $\hat{\eta}$, defined by

$$\hat{\eta}(t) = \begin{cases} 90.5, & -0.534 < t \leq 0, \\ 0.167, & -0.66 < t \leq -0.534, \\ 1.0, & t \leq -0.66, \end{cases} \quad (8)$$

which was estimated from simulated data (Appendix B.2)

Overall, we found that using the constant table $\mathcal{L}_{\text{const}}$ leads to very noisy and biased estimates of ρ (as might be expected from Figure 4). The other tables $\mathcal{L}_{\text{exact}}$, $\mathcal{L}_{\text{approx}}$, and $\mathcal{L}_{\text{miss}}$ all lead to much more accurate estimates. Using $\mathcal{L}_{\text{exact}}$ (the exact likelihood table with the true size history) produces slightly more accurate estimates of ρ than using $\mathcal{L}_{\text{approx}}$ or $\mathcal{L}_{\text{miss}}$. However, the three non-constant tables $\mathcal{L}_{\text{exact}}$, $\mathcal{L}_{\text{approx}}$, and $\mathcal{L}_{\text{miss}}$ all produce very similar results that are hard to distinguish from one another.

Figures 5 and 6 show the accuracy of estimated recombination maps $\hat{\rho}$ on simulated data. We simulated $n = 20$ sequences under the 3-epoch demography defined in (6), with the true maps ρ taken from previous estimates for the X chromosome of *D. melanogaster* (Chan et al., 2012). In all, there were 110 independent datasets, with estimated maps $\hat{\rho}$ of length 500 kb. See Appendix B for further details.

In Figure 5, we plot ρ and $\hat{\rho}$ for a particular 500 kb region. Qualitatively, the constant size estimate $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ is less accurate and has wilder fluctuations. Figure 6 shows that over all 110 replicates, the constant size estimate $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ has high bias and low correlation with the truth, compared to the estimates $\hat{\rho}_{\mathcal{L}_{\text{exact}}}$, $\hat{\rho}_{\mathcal{L}_{\text{miss}}}$, and $\hat{\rho}_{\mathcal{L}_{\text{approx}}}$ which account for variable demography. Following Wegmann et al. (2011), we plot the correlation of $\hat{\rho}$ with ρ at multiple scales; at all scales, the constant-demography estimate $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ is considerably worse than the other estimates. In general, using an inferred demography or an approximate lookup table results in only a very mild reduction in accuracy compared to using the true sampling probabilities.

We also considered a constant (flat) map ρ ; the estimated $\hat{\rho}$ are shown in Figure 7. Consistent with Johnston and Cutler (2012), we find that the constant-demography estimate $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ can have extreme peaks, and is generally very noisy. On the other hand, the estimates $\hat{\rho}_{\mathcal{L}_{\text{exact}}}$, $\hat{\rho}_{\mathcal{L}_{\text{miss}}}$, $\hat{\rho}_{\mathcal{L}_{\text{approx}}}$ that account for demography have less noise and fewer large deviations.

5 Runtime and Accuracy of Likelihoods

5.1 Runtime of the exact and approximate likelihood formulas

Both the approximate (2) and exact likelihood formulas (Theorem 1) require computing products $\mathbf{v}e^{\mathbf{A}}$ for some $k \times k$ matrix \mathbf{A} and $1 \times k$ row vector \mathbf{v} . Naively, this kind of vector-matrix multiplication costs $O(k^2)$. However, in our case \mathbf{A} is sparse, with $O(k)$ nonzero entries, allowing us to compute $\mathbf{v}e^{\mathbf{A}}$ up to numerical precision in $O(k\mathcal{T})$ time, where \mathcal{T} is some finite number of sparse matrix-vector products depending on \mathbf{A} (Al-Mohy and Higham, 2011). In particular, $k = O(n^3)$ for the approximate formula, whereas $k = O(n^6)$ for the exact formula. Thus, computing the likelihood table costs $O(n^3\mathcal{T})$ and $O(n^6\mathcal{T})$ for the approximate (2) and exact (Theorem 1) formulas, respectively. See Appendix C for a more detailed analysis of the computational complexity, and a description of the algorithm for computing $\mathbf{v}e^{\mathbf{A}}$. Note that \mathcal{T} depends nontrivially on the sample size n , as well as the parameters ρ , θ , η_d , and t_d .

We present running times for LDpop in Figure 8. We used LDpop to generate likelihood tables with $\rho \in \{0, 1, \dots, 100\}$, using 24 cores on a computer with 256 GB of RAM. On the 3-epoch






Figure 5: Comparison of recombination maps inferred using different lookup tables. We simulated $n = 20$ haplotypes under the 3-epoch model (6) and using the recombination map shown in black dashed line. (a) Inferred map $\hat{\rho}_{\mathcal{L}_{\text{exact}}}$ obtained using the exact likelihoods for the true demography. (b) Inferred map $\hat{\rho}_{\mathcal{L}_{\text{miss}}}$ obtained using the empirically estimated demography in (8). (c) Inferred map $\hat{\rho}_{\mathcal{L}_{\text{approx}}}$ obtained using an approximate lookup table for the true demography. (d) Inferred map $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ obtained assuming a constant population size of $\eta \equiv 1$. Note that $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ is much noisier than the other estimates, while using an inferred demography or an approximate lookup table results in only a very mild reduction in accuracy compared to using the true sampling probabilities. These $\hat{\rho}$ were produced with LDhelmet; using LDhat led to very similar results.



(a)



(b)

Figure 6: Accuracy of the estimated maps $\hat{\rho}_{\mathcal{L}_{\text{exact}}}$, $\hat{\rho}_{\mathcal{L}_{\text{miss}}}$, $\hat{\rho}_{\mathcal{L}_{\text{approx}}}$, $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ over 110 simulations similar to Figure 5. The estimate $\hat{\rho}_{\mathcal{L}_{\text{const}}}$ obtained assuming constant demography is substantially more biased and noisier than the other estimates. (a) The average per-base bias $\hat{\rho} - \rho$. (b) The square Pearson correlation coefficient R^2 over different scales. This R^2 is distinct from the r^2 statistic measuring linkage disequilibrium (Section 4.1). To compute R^2 for scale s , the middle 500 kb region of each 1 Mb simulation was divided into non-overlapping windows of size s and we compared the average of $\hat{\rho}$ to the average of ρ in each window. The x -axis is stretched by $x \mapsto \sqrt{x}$.

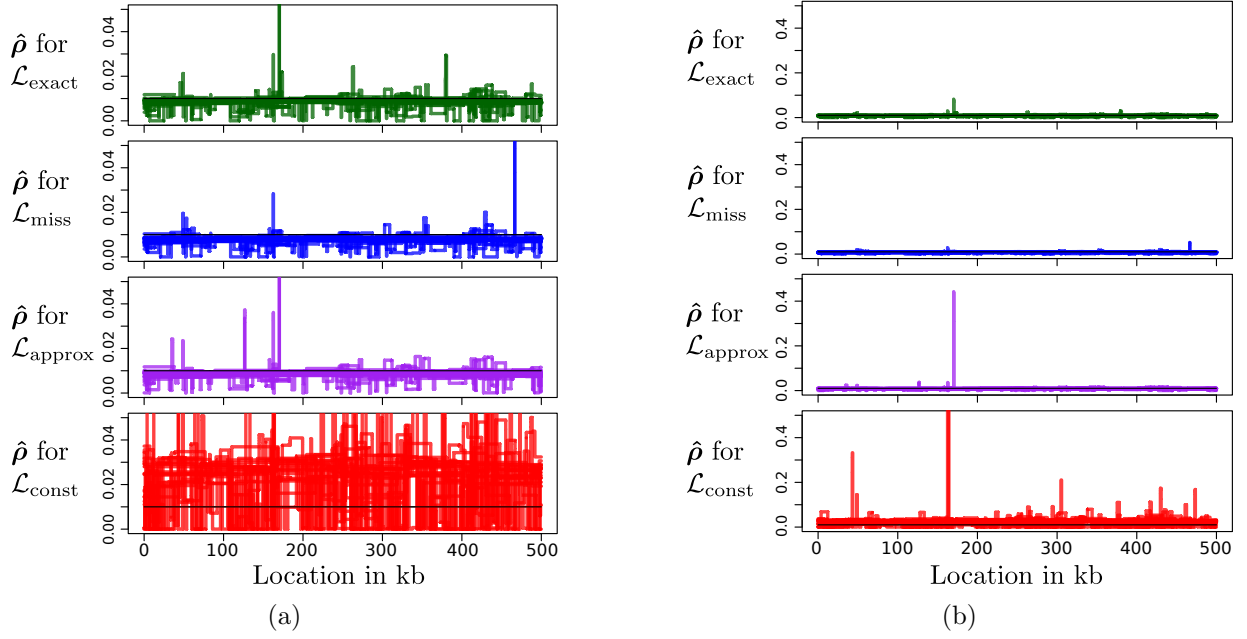




Figure 7: Inferred recombination maps $\hat{\rho}$ produced by LDhelmet using four different look up tables ($\mathcal{L}_{\text{const}}$, $\mathcal{L}_{\text{exact}}$, $\mathcal{L}_{\text{approx}}$, $\mathcal{L}_{\text{miss}}$), when the true recombination rate is constant at $\rho = 0.01$ per bp. Results for 20 simulated datasets are shown. Each simulation was done under the 3-epoch demography defined in (6). Results for LDhat were very similar (not shown). (a) Inferred recombination maps $\hat{\rho}$. (b) The same plot but zoomed out. For the figure corresponding to $\mathcal{L}_{\text{const}}$, the highest peak reaches $\rho = 2.5$ (not displayed), which is 250 times the true value.



(a)



(b)

Figure 8: Runtime of LDpop to compute a likelihood table with $\rho \in \{0, 1, \dots, 100\}$. Experiments were performed using 24 cores on a computer with 256 GB of RAM. (a) Runtime as a function of sample size, for both the approximate and exact formulas, and for two demographies (a constant size history and the 3-epoch model (6)). (b) Runtime as a function of the number of epochs \mathcal{D} on $[-T_{\max}, 0]$, for the exact formula with $n = 20$, with 10 repetitions with random population sizes. Notice that runtime depends more on T_{\max} than \mathcal{D} , and grows sublinearly with \mathcal{D} (a 2-fold increase of \mathcal{D} yields less than a 2-fold increase in runtime).




Figure 9: The approximate table $\{\log \hat{\mathbb{P}}(\mathbf{n})\}$ plotted against the exact table $\{\log \mathbb{P}(\mathbf{n})\}$, for a lookup table with $n = 20$ and $\rho \in \{0, 1, \dots, 100\}$, under the 3-epoch model in (6). (a) $N = 20$ Moran particles. (b) $N = 100$ Moran particles. Note that the approximate table with $N = 100$ is extremely accurate, and visually indistinguishable from the true values.

demography (6), we ran the exact formula up to $n = 40$ (17 hours), and the approximate formula up to $n = 256$ (13 hours). The constant demography takes nearly the same amount of time as the 3-epoch demography, which agrees with our general experience that computing the initial stationary distribution $\tilde{\mathbf{\Lambda}}^{-D}$ is more expensive than multiplying the matrix exponentials $e^{\tilde{\mathbf{\Lambda}}^{dt}}$. Using faster algorithms to compute $\tilde{\mathbf{\Lambda}}^{-D}$ should lead to substantial improvements in runtime.

Figure 8 also examines how LDpop scales with the number of epochs in the demography. We split $[-T_{\max}, 0]$ into \mathcal{D} intervals of length $\frac{T_{\max}}{\mathcal{D}}$, each with a random population size $\frac{1}{\eta_d} \sim \text{logUniform}(0.1, 10)$, with 10 repetitions per setting of T_{\max} and \mathcal{D} . Empirically, LDpop scales sublinearly with \mathcal{D} , and LDpop has no problem handling $\mathcal{D} = 64$ epochs. We also note that T_{\max} has a greater impact on runtime than \mathcal{D} ; this is because the matrix exponentials are essentially computed by solving an ODE from $-T_{\max}$ to 0, as noted in Appendix C.1.

5.2 Accuracy of the approximate likelihood

In Section 4.2 we found that using the approximate likelihood (2) has little impact on recombination rate estimation, suggesting that it is an accurate approximation to the exact formula in Theorem 1.

We examine this in greater detail in Figure 9, for $n = 20$, the 3-epoch demography (6), and a lookup table with $\rho \in \{0, 1, \dots, 100\}$. We compare the approximate against the exact values for $N = 20$ and $N = 100$ Moran particles in the approximate model. The approximate table with $N = 20$ is reasonably accurate, with some mild deviations from the truth. The approximate table with $N = 100$ is extremely accurate, and visually indistinguishable from the true values.

5.3 Runtime and accuracy of the importance sampler

We plot the runtime of the importance sampler (Section 3.2) in Figure 10a. For the previous 3-epoch demography (6), we drew $K = 200$ importance samples for each of the 275 configurations \mathbf{n} with $n = n^{(c)} = 20$, with $\rho \in \{0.1, 1, 10, 100\}$. The runtime of the importance sampler generally

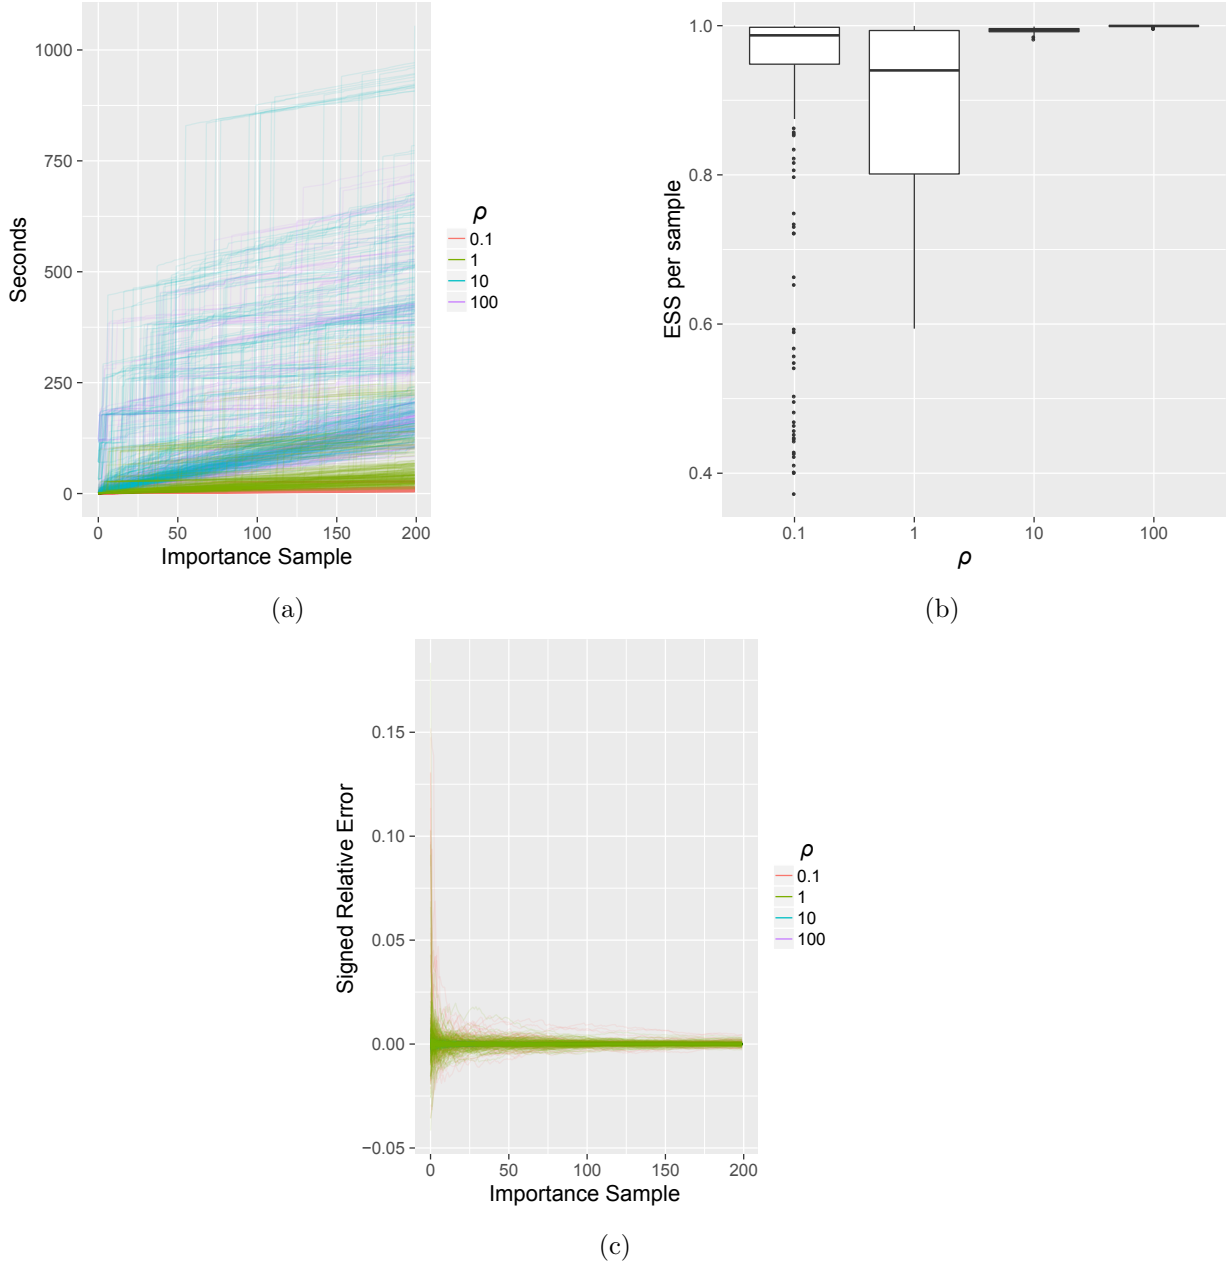


Figure 10: Accuracy and runtime of importance sampling, on the 3-epoch demography (6) with $\theta = 0.008$ and $\rho = 1.0$, drawing 200 genealogies for each of the 275 fully-specified configurations \mathbf{n} with $n = 20$. (a) Runtime for each \mathbf{n} , as a function of the number of importance samples. Higher ρ generally took a longer time. Using 20 cores, the time to sample all 275 configurations took about 4 minutes when $\rho = 0.1$, but 1 hour for $\rho = 100$. (b) The ESS per importance sample, for each configuration \mathbf{n} . (c) The signed relative error $\frac{\text{Est} - \text{Truth}}{\text{Truth}}$ of $\log \hat{\mathbb{P}}(\mathbf{n})$, as a function of the number of importance samples. The true values were computed via Theorem 1.

increased with ρ ; using 20 cores, sampling all 275 configurations took about 4 minutes with $\rho = 0.1$, but about 1 hour with $\rho = 100$. We further analyze the computational complexity of the importance sampler in Appendix C.4.

The number K of importance samples required to reach a desired level of accuracy is typically measured with the effective sample size (ESS):

$$\text{ESS} = \frac{\left(\sum_{k=1}^K w_k\right)^2}{\sum_{k=1}^K w_k^2},$$

where $w_k = \frac{d\mathbb{P}(\mathbf{n}_{\leq 0}^{(k)})}{d\hat{Q}(\mathbf{n}_{\leq 0}^{(k)})}$ denotes the importance weight of the k th sample (see (4)). Note that $\text{ESS} \leq K$ always, with equality only achieved if the w_k have 0 variance.

Compared to previous coalescent importance samplers, our proposal distribution is highly efficient. We plot the ESS per importance sample (i.e. $\frac{\text{ESS}}{K}$) in Figure 10b. Typically $\frac{\text{ESS}}{K} > 0.8$; for $\rho \in \{10, 100\}$, the ESS is close to its optimal value, with $\frac{\text{ESS}}{K} \approx 1$. In Figure 10c, we compare the log-likelihood estimated from importance sampling to the true value computed with Theorem 1; after $K = 200$ importance samples, the signed relative error is well under 1% for all \mathbf{n} .

By contrast, the previous two-locus importance sampler of Fearnhead and Donnelly (2001), which assumes a constant population size, achieves ESS anywhere between $0.05K$ and $0.5K$, depending on \mathbf{n}, θ, ρ (result not shown). This importance sampler is based on a similar result as Theorem 2, with optimal rates $\phi_{\mathbf{n}, \mathbf{m}} \frac{\mathbb{P}(\mathbf{m})}{\mathbb{P}(\mathbf{n})}$. However, to approximate $\frac{\mathbb{P}(\mathbf{m})}{\mathbb{P}(\mathbf{n})}$, previous approaches did not use a Moran model, but followed the approach of Stephens and Donnelly (2000), using an approximate “conditional sampling distribution” (CSD). We initially tried using the CSD of Fearnhead and Donnelly (2001) and later generalizations to variable demography (Sheehan et al., 2013; Steinrücken et al., 2015), but found that importance sampling failed under population bottleneck scenarios, with the ESS repeatedly crashing to lower and lower values. Previous attempts to perform importance sampling under variable demography (Ye et al., 2013) have also encountered low ESS, though in the context of an infinite sites model (as opposed to a 2-locus model). However, Dialdestoro et al. (2016) recently devised an efficient 2-locus importance sampler using the CSD-approach, in conjunction with advanced importance sampling techniques. Their importance sampler allows archaic samples and is thus time-inhomogeneous, but it only models constant population size histories.

6 Discussion

In this paper, we have developed a novel algorithm for computing the exact two-locus sampling probability under arbitrary piecewise-constant demographic histories. These two-locus likelihoods can be used to study the impact of demography on LD and also to improve fine-scale recombination rate estimation. Indeed, using two-locus sampling probabilities computed under the true or an inferred demography, we were able to obtain recombination rate estimates with substantially less noise and fewer spurious peaks that could potentially be mistaken for hotspots.

We have implemented our method in a freely available software package, LDpop. This program also includes an efficient approximation to the true sampling probability that easily scales to hundreds in sample size. In practice, highly accurate approximations to the true sampling probability for sample size n can be obtained quickly by first applying the approximate algorithm with N Moran particles larger than n and then down-sampling to the desired sample size n .

In principle, one could also obtain an accurate approximation to the sampling probability using

our importance sampler, which is also implemented in LDpop. We have not optimized this code, however, and we believe that its main utility will be in sampling two-locus ARGs from the posterior distribution. Lastly, we note that, in addition to improving the inference of fine-scale recombination rate variation, our two-locus likelihoods can be utilized in other applications such as hotspot hypothesis testing and demographic inference.

Acknowledgments

This research is supported in part by an NIH grant R01-GM108805, an NIH training grant T32-HG000047, and a Packard Fellowship for Science and Engineering.

Appendix

A Proposal distribution of importance sampler

For the importance sampler of Section 3.2, we construct the proposal distribution $\hat{Q}(\mathbf{n}_{\leq 0})$ by approximating the optimal proposal distribution $Q_{\text{opt}}(\mathbf{n}_{\leq 0}) = \mathbb{P}(\mathbf{n}_{\leq 0} \mid \mathbf{n}_0)$ given in Theorem 2. We start by choosing a grid of points $-\infty < \tau_1 < \tau_2 < \dots < \tau_J = 0$, then set \hat{Q} to be a backwards in time Markov chain, whose rates at $t \in (\tau_j, \tau_{j+1})$ are the linear interpolation

$$\hat{q}_{\mathbf{n}, \mathbf{m}}^{(t)} = \frac{\tau_{j+1} - t}{\tau_{j+1} - \tau_j} \hat{q}_{\mathbf{n}, \mathbf{m}}^{(\tau_j)} + \frac{t - \tau_j}{\tau_{j+1} - \tau_j} \hat{q}_{\mathbf{n}, \mathbf{m}}^{(\tau_{j+1})}, \quad (9)$$

with the rates at the grid points given by

$$\hat{q}_{\mathbf{n}, \mathbf{m}}^{(\tau_j)} = \begin{cases} \phi_{\mathbf{n}, \mathbf{m}}^{(\tau_j)} \frac{\hat{\mathbb{P}}_{\tau_j}(\mathbf{m})}{\hat{\mathbb{P}}_{\tau_j}(\mathbf{n})}, & \text{if } \mathbf{m} \neq \mathbf{n}, \\ -\sum_{\nu \neq \mathbf{n}} \hat{q}_{\mathbf{n}, \nu}^{(\tau_j)}, & \text{if } \mathbf{m} = \mathbf{n}, \end{cases}$$

with $\hat{\mathbb{P}}_{\tau_j}(\mathbf{n})$ an approximation to the likelihood $\mathbb{P}_{\tau_j}(\mathbf{n})$. In particular, we set $\hat{\mathbb{P}}_{\tau_j}(\mathbf{n}) = \mathbb{P}_{\tau_j}^{(N)}(\mathbf{n})$, using the approximate likelihood formula (2) in Section 2.3. The approximate likelihoods $\{\mathbb{P}_{\tau_j}^{(N)}(\mathbf{n})\}$ can be efficiently computed along a grid of points using the method of Appendix C.1 (see also Appendix C.4).

To sample from \hat{Q} , we note that for configuration \mathbf{n} at time t , the time $S < t$ of the next event has CDF $\mathbb{P}(S < s) = \exp(-\int_s^t \hat{q}_{\mathbf{n}, \mathbf{n}}^{(u)} du)$ for $s < t$. Thus, S can be sampled by first sampling $X \sim \text{Uniform}(0, 1)$, and then solving for $\log(X) = -\int_S^t \hat{q}_{\mathbf{n}, \mathbf{n}}^{(u)} du$ via the quadratic formula (since $\hat{q}_{\mathbf{n}, \mathbf{n}}^{(u)}$ is piecewise linear; see (9)). Having sampled S , we can then sample the next configuration \mathbf{m} with probability $-\hat{q}_{\mathbf{n}, \mathbf{m}}^{(S)} / \hat{q}_{\mathbf{n}, \mathbf{n}}^{(S)}$.

B Details of simulation study in Section 4.2

B.1 Simulated data

We simulated independent 1 Mb segments with $n = 20$ haplotypes under the 3-epoch demography $\eta(t)$ in (6). To do so, we generated trees using the program MaCS (Chen et al., 2009), and then generated mutations according to a quadra-allelic mutational model. For the variable recombination maps used in Figures 5 and 6, we divided the recombination map of the X chromosome of *D. melanogaster* from Raleigh, NC inferred by Chan et al. (2012) into 22 non-overlapping 1 Mb windows and simulated 5 replicates, for a total of 110 Mb of simulated data. For the constant map used in Figure 7, we generated 20 datasets with $\rho = 0.01$ per base.

B.2 Estimation of misspecified demography $\hat{\eta}(t)$

To estimate the misspecified demography $\hat{\eta}(t)$ of (8), we pooled all bi-allelic SNPs from the 110 simulated segments of the variable recombination map, and then used the folded site frequency spectrum (SFS) of the simulated SNPs to estimate $\hat{\eta}(t)$. Specifically, we fit $\hat{\eta}(t)$ by maximizing a composite likelihood, viewing each SNP as an independent draw from a multinomial distribution proportional to the expected SFS. We computed the expected SFS with the software package

`moni` (Kamm et al., 2016), and fixed the most ancient population size to 1.0 due to scaling and identifiability issues.

B.3 Recombination map estimation

After removing all non-biallelic SNPs, we ran both LDhat and LDhelmet on the resulting data, using a block penalty of 50 as recommended by Chan et al. (2012) for *Drosophila*-like data (the block penalty is a tuning parameter that is multiplied by the number of change-points in the estimated map $\hat{\rho}$, and added to the log composite likelihood; thus a high block penalty discourages over-fitting). We took the posterior median inferred at each position to be the estimated map $\hat{\rho}$. We only used the centermost 500 kb of each estimate $\hat{\rho}$ to avoid the issue of edge effects.

C Computational complexity

C.1 Computing the action of a sparse matrix exponential

Both Theorem 1 and the approximate formula (2) rely on “the action of the matrix exponential” (Al-Mohy and Higham, 2011). Let \mathbf{A} be a $k \times k$ matrix and \mathbf{v} a $1 \times k$ row vector. We need to compute expressions of the form $\mathbf{v}e^{\mathbf{A}}$. Naively, this kind of vector-matrix multiplication costs $O(k^2)$. However, in our case \mathbf{A} will be sparse, with k nonzero entries, allowing us to more efficiently compute $\mathbf{v}e^{\mathbf{A}}$.

In particular, we use the algorithm of Al-Mohy and Higham (2011), as implemented in the Python package `scipy`. For $s \in \mathbb{Z}_+$, define $T_m(s^{-1}\mathbf{A}) = \sum_{i=0}^m (s^{-1}\mathbf{A})^i/i!$, the truncated Taylor series approximation of $e^{s^{-1}\mathbf{A}}$. Then, we have

$$\mathbf{v}e^{\mathbf{A}} = \mathbf{v} \left(e^{s^{-1}\mathbf{A}} \right)^s \approx \mathbf{v}[T_m(s^{-1}\mathbf{A})]^s.$$

Now let $\mathbf{b}_j = \mathbf{v}[T_m(s^{-1}\mathbf{A})]^j$, so \mathbf{b}_j is a $1 \times k$ row vector. Then

$$\mathbf{b}_j = \mathbf{b}_{j-1}T_m(s^{-1}\mathbf{A}) = \sum_{i=0}^m \mathbf{b}_{j-1} \frac{(s^{-1}\mathbf{A})^i}{i!},$$

with $\mathbf{v}e^{\mathbf{A}} \approx \mathbf{b}_s$, and \mathbf{b}_s evaluated in $\mathcal{T} = ms$ vector-matrix multiplications, each costing $O(k)$ by the sparsity of \mathbf{A} . Approximating $\mathbf{v}e^{\mathbf{A}}$ thus costs $O(\mathcal{T}k)$ time. Both m, s are chosen automatically to bound

$$\frac{\|\Delta\mathbf{A}\|_1}{\|\mathbf{A}\|_1} \leq \text{tolerance} \approx 1.1 \times 10^{-16},$$

with $\Delta\mathbf{A}$ defined by $[T_m(s^{-1}\mathbf{A})]^s = e^{\mathbf{A} + \Delta\mathbf{A}}$, and the matrix norm given by $\|\mathbf{A}\|_1 = \sup_{\mathbf{w} \neq \mathbf{0}} \frac{\|\mathbf{w}\mathbf{A}\|_1}{\|\mathbf{w}\|_1}$. To avoid numerical instability, m is also bounded by $m \leq m_{\max} = 55$. Al-Mohy and Higham (2011) provide some analysis for the size of $\mathcal{T} = ms$, but this analysis is rather involved. Very roughly, \mathcal{T} is proportional to $\|\mathbf{A}\|$ (for arbitrary matrix norm $\|\cdot\|$), so computing $\mathbf{v}e^{2\mathbf{A}}$ takes twice as long as $\mathbf{v}e^{\mathbf{A}}$, and computing $\mathbf{v}e^{t\mathbf{A}}$ is roughly proportional to t . This is because $\mathbf{v}e^{t\mathbf{A}}$ is essentially computed by numerically integrating the ODE $\nabla \mathbf{f}(s) = \mathbf{f}(s)\mathbf{A}$ for $s \in [0, t]$.

We note that $\mathbf{b}_j \approx \mathbf{v}e^{s^{-1}j\mathbf{A}}$, and thus this algorithm approximates $\mathbf{v}e^{t\mathbf{A}}$ along a grid of points $t \in \{s^{-1}, 2s^{-1}, \dots, 1\}$. If $\mathbf{v}e^{t\mathbf{A}}$ is needed at additional points, then extra grid points can be added at those times.

C.2 Complexity of the exact likelihood formula (Theorem 1)

We consider the computational complexity of computing $\mathbb{P}_0(\mathbf{n})$ via Theorem 1. Note that the formula (3) simultaneously computes $\mathbb{P}_0(\mathbf{n})$ for all configurations $\mathbf{n} \in \mathcal{N}$.

As usual, we assume 2 alleles $\mathcal{A} = \{0, 1\}$, as is assumed by LDhat and the applications in Section 4. We start by considering the dimensions of the sampling probability vectors \mathbf{p}^d and rate matrices $\tilde{\mathbf{\Lambda}}^d$ for intervals $(t_d, t_{d+1}]$. The set of a, b, c haplotypes is $H = \{00, 01, 10, 11, 0*, 1*, *0, *1\}$, so $|H| = 8$. Thus, there are $O(n^6)$ possible configurations \mathbf{n} with $n^{(a)} = n^{(b)} = n - n^{(c)}$. In particular, there are $O(n^6)$ ways to specify $n_{00}, n_{01}, n_{10}, n_{11}, n_{0*}, n_{*0}$, and then $n_{1*} = n - \sum_{i,j \in \{0,1\}} n_{ij} - n_{0*}$ and $n_{*1} = n - \sum_{i,j \in \{0,1\}} n_{ij} - n_{*0}$ are determined. Thus, \mathbf{p}^d is a row vector of dimension $1 \times O(n^6)$, and $\tilde{\mathbf{\Lambda}}^d$ is a square matrix of dimension $O(n^6) \times O(n^6)$, but $\tilde{\mathbf{\Lambda}}^d$ is sparse, with only $O(n^6)$ nonzero entries.

By using the algorithm of Appendix C.1, we can compute $\mathbf{p}^{d+1} = [(\mathbf{p}^d \odot \tilde{\gamma}^d) e^{\tilde{\mathbf{\Lambda}}^d(t_{d+1}-t_d)}] \div \tilde{\gamma}^d$ from \mathbf{p}^d in $O(\mathcal{T}_d n^6)$ time, where \mathcal{T}_d is the number of vector-matrix multiplications to compute the action of $e^{\tilde{\mathbf{\Lambda}}^d(t_{d+1}-t_d)}$. We note that the stationary distribution $(\gamma_0^d, \dots, \gamma_n^d)$ can be computed in $n+1$ steps: $\mathbf{\Gamma}^d$ is the rate matrix of a simple random walk with $n+1$ states, so $\gamma_{i+1}^d = \gamma_i^d \frac{[\mathbf{\Gamma}^d]_{i,i+1}}{[\mathbf{\Gamma}^d]_{i+1,i}}$ and $\sum_i \gamma_i^d = 1$.

Similarly, the initial value $\mathbf{p}^{-D+1} = \tilde{\mathbf{\Lambda}}^{-D} \div \tilde{\gamma}^{-D}$ can be computed via sparse vector-matrix multiplications, using the technique of power iteration. For $\mu = \frac{1}{\max_{ij} [\tilde{\mathbf{\Lambda}}^{-D}]_{ij}}$ and arbitrary positive vector $\mathbf{v}^{(0)}$ with $\|\mathbf{v}^{(0)}\|_1 = 1$, we have $\mathbf{v}^{(i)} := \mathbf{v}^{(0)}(\mu \tilde{\mathbf{\Lambda}}^{-D} + I)^i \rightarrow \tilde{\mathbf{\Lambda}}^{-D}$ as $i \rightarrow \infty$. In particular, we set the number of iterations, \mathcal{T}_{-D} , so that $\|\log(\mathbf{v}^{(\mathcal{T}_{-D})}) \div \mathbf{v}^{(\mathcal{T}_{-D}-1)}\|_1 < 1 \times 10^{-8}$, where $\log \mathbf{v}$ is the element-wise log of \mathbf{v} . As noted in Section 5.1, in practice we found $\mathcal{T}_{-D} \gg \mathcal{T}_d$ for $d > -D$, i.e. computing the initial stationary distribution $\tilde{\mathbf{\Lambda}}^{-D}$ was more expensive than multiplying the matrix exponentials $e^{\tilde{\mathbf{\Lambda}}^d t}$.

To summarize, computing $\mathbb{P}_0(\mathbf{n})$ for all $O(n^6)$ configurations $\mathbf{n} \in \mathcal{N}$ of size n costs $O(n^6 \mathcal{T})$, with $\mathcal{T} = \sum_{d=-D}^{-1} \mathcal{T}_d$. We caution that \mathcal{T} depends on $n, \{t_d\}, \{\tilde{\mathbf{\Lambda}}^d\}$.

The memory cost of Theorem 1 is $O(n^6)$, since $\tilde{\mathbf{\Lambda}}^d$ has $O(n^6)$ nonzero entries.

C.2.1 Comparison with Golding's equations

Under constant population size, Golding (1984) proposed a method to compute $\mathbb{P}_0(\mathbf{n})$ by solving a linear system of equations $\mathbf{g}\mathbf{G} = \mathbf{g}$, where $\mathbf{g} = [\mathbb{P}_0(\mathbf{n})]_{\mathbf{n} \in \mathcal{N}'}$ is the vector of sampling probabilities indexed by the $O(n^8)$ configurations $\mathcal{N}' = \{\mathbf{n} : \max(n^{(a)} + n^{(c)}, n^{(b)} + n^{(c)}) \leq n\}$ with at most n alleles at each locus. Hudson (2001) solves this linear system, costing $O(n^8 \mathcal{T})$ where (as above) \mathcal{T} is some finite number of sparse matrix-vector multiplications.

For the case of constant population size, Theorem 1 reduces to solving a sparse system $\tilde{\mathbf{\Lambda}}^{-D} \tilde{\mathbf{\Lambda}}^{-D} = \tilde{\mathbf{\Lambda}}^{-D}$, which is similar in spirit to solving Golding's equations $\mathbf{g}\mathbf{G} = \mathbf{g}$. The $O(n^6 \mathcal{T})$ runtime of Theorem 1 at first seems superior to the $O(n^8 \mathcal{T})$ runtime of Golding's equations, but in fact the number of matrix multiplications \mathcal{T} is not comparable between the two methods. Most importantly, Hudson (2001) exploits the structure of the $O(n^8)$ equations to decompose them into smaller sub-systems of $O(n^4)$ equations, which may lead to smaller \mathcal{T} . Algorithmic details also lead to important differences: we use power iteration, whereas Hudson (2001) use a conjugate gradient method, with less stringent convergence criteria (stopping when the relative l_2 error for each sub-system of equations is $< 10^{-4}$).

Preliminary tests suggest that the C code of Hudson (2001), and a similar implementation by Chan et al. (2012), are faster than our current method for solving $\tilde{\mathbf{\Lambda}}^{-D} \tilde{\mathbf{\Lambda}}^{-D} = \tilde{\mathbf{\Lambda}}^{-D}$. We are planning future updates to LDpop that will speed up the initial stationary distribution $\tilde{\mathbf{\Lambda}}^{-D}$, either

by changing the algorithmic details of our solver, or by using Golding’s equations to compute $\tilde{\lambda}^{-D}$ instead.

C.3 Complexity of approximate likelihood formula

The method of computing the approximate likelihood formula (2) is similar to computing Theorem 1, in that we can compute an initial stationary distribution $\lambda_{(N)}^{-D}$ by power iteration, and then propagate it forward in time by applying the action of the sparse matrix exponential $e^{\Lambda_{(N)}^d(t_{d+1}-t_d)}$. However, instead of $O(n^6)$ states, there are $O(N^3)$ total states: there are 4 possible fully-specified haplotypes $\{00, 01, 10, 11\}$, and thus the requirement that the number of lineages sums to N yields $O(N^3)$ possible states for the Moran model \mathbf{M}_t . Thus, computing the approximate likelihood formula (2) costs $O(N^3\mathcal{T})$ time and $O(N^3)$ memory space.

C.4 Complexity of importance sampler

Here we examine the computational complexity of the importance sampler of Section 3.2.

To construct the proposal distribution \hat{Q} (Appendix A), we must compute approximate likelihoods $\mathbb{P}_t^{(N)}(\mathbf{n})$ defined by (2) along a grid of points $t \in \{\tau_1, \tau_2, \dots, \tau_J\}$. We start by computing the Moran likelihoods $\{\mathbb{P}_{(N)}(\mathbf{M}_{\tau_j})\}_j$ using the action of the sparse matrix exponential. As discussed in Appendix C.1, the method of Al-Mohy and Higham (2011) yields $\{\mathbb{P}_{(N)}(\mathbf{M}_{\tau_j})\}_j$ as a byproduct of computing $\mathbb{P}_{(N)}(\mathbf{M}_0)$ at essentially no extra cost. Thus, computing the terms $\{\mathbb{P}_{(N)}(\mathbf{M}_t)\}$ costs $O(N^3\mathcal{T})$ (absorbing the minor cost of an additional J extra grid points into \mathcal{T}).

We then compute $\mathbb{P}_t^{(N)}(\mathbf{n})$ by subsampling from $\mathbb{P}_{(N)}(\mathbf{M}_t)$ as in (2), and thus set $N = 2n$, since $2n$ is the maximum number of individuals in \mathbf{n} (because each of the original n lineages can recombine into two lineages). However, it is inefficient to compute $\hat{\mathbb{P}}_t(\mathbf{n})$ by subsampling for every value of \mathbf{n} separately. Instead, it is better to use a dynamic program $\hat{\mathbb{P}}_t(\mathbf{n}) = \sum_{\mathbf{m}} \hat{\mathbb{P}}_t(\mathbf{m})\mathbb{P}(\mathbf{n} | \mathbf{m})$, where the sum is over all configurations \mathbf{m} obtained by adding an additional sample to \mathbf{n} .

This costs $O(n^8J)$ time and space, since there are J grid points and $O(n^8)$ possible configurations of \mathbf{n} . Then, assuming a reasonably efficient proposal, the expected cost to draw K importance samples is $O(nJK)$, since the expected number of coalescence, mutation, and recombination events before reaching the marginal common ancestor at each locus is $O(n)$ (Griffiths, 1991). This approach thus takes $O(n^3\mathcal{T} + n^8J + n^4JK)$ expected time to compute $\mathbb{P}_0(\mathbf{n})$ for all $O(n^3)$ possible \mathbf{n} . In practice, we only precomputed $\hat{\mathbb{P}}_t(\mathbf{n})$ for the $O(n^4)$ fully specified \mathbf{n} (without missing alleles), but computed and cached $\hat{\mathbb{P}}_t(\mathbf{n})$ as needed for partially specified \mathbf{n} (with missing alleles). The theoretical running time to compute the full lookup table is still $O(n^3\mathcal{T}D + n^8J + n^4JK)$, but in practice, many values of \mathbf{n} are highly unlikely and never encountered at each τ_j .

D Proofs

For a stochastic process $\{X_t\}_{t \leq 0}$, we denote its partial sample paths with the following notation: $X_{s:t} = \{X_u : u \in (s, t]\}$ and $X_{\leq s} = X_{-\infty:s}$.

D.1 Proof of Theorem 1

We start by constructing a *forward-in-time* Markov jump process $\tilde{\mathbf{M}}_{\leq 0}$ with state space $\mathcal{N} = \{\mathbf{n} : n^{(abc)} = (k, k, n - k), 0 \leq k \leq n\}$. $\tilde{\mathbf{M}}_t$ changes due to four types of events: mutation, copying, “recoalescence”, and “recombination”.

1. Individual alleles mutate at rate $\frac{\theta}{2}$ according to transition matrix \mathbf{P} .
2. Lineages copy their alleles onto each other, with the rate depending on the lineage type. Each pair of a types experiences a copying event at rate $\frac{1}{\eta_d}$, with the direction of copying chosen with probability $\frac{1}{2}$. The rates are the same for every pair of b and every pair of c types. Pairs of (a, c) and (b, c) types also experience copying at rate $\frac{1}{\eta_d}$, however the direction of copying is always from the c type to the a or b type, and only happens at one allele (left for a , right for b).
3. a types do not copy onto b types, and vice versa. Instead, they merge (“recoalesce”) into a single c type at rate $\frac{1}{\eta_d}$ per pair. Note this is similar to the coalescent, but here the recoalescence happens *forward in time* rather than *backwards in time*.
4. Each c type splits into a and b types at rate $\frac{\rho}{2}$. Again, this is similar to the coalescent, however here the “recombination” happens *forward in time*, while in the coalescent with recombination it happens at rate $\frac{\rho}{2}$ going *backwards in time*.

Then $\tilde{\mathbf{M}}_t$ has forward-in-time rate matrix $\mathbf{\Lambda}^d$ in $(t_d, t_{d+1}]$, with $\mathbf{\Lambda}^d$ given in Table 2.

Now let C_t denote the number of c types in $\tilde{\mathbf{M}}_t$ (so the number of a and b types are each $n - C_t$). Note that C_t is unaffected by mutation and copying events, and so C_{t+h} is conditionally independent of $\tilde{\mathbf{M}}_t$ given C_t , for $h \geq 0$. Thus C_t is a Markov jump process with rate matrix $\mathbf{\Gamma}^d$ in $(t_d, t_{d+1}]$, where $\mathbf{\Gamma}^d$ is a tridiagonal square matrix indexed by $\{0, 1, \dots, n\}$, with $\Gamma_{m, m-1}^d = \frac{\rho}{2}m$, $\Gamma_{m, m+1}^d = (n - m)^2 \frac{1}{\eta_d}$, and $\Gamma_{m, m}^d = -\Gamma_{m, m-1}^d - \Gamma_{m, m+1}^d$.

We can therefore sample $\tilde{\mathbf{M}}_{\leq 0}$ in two steps, as illustrated in Figure 3:

1. First, sample the recoalescence and recombination events. In other words, sample $C_{\leq 0}$ using its rate matrices $\{\mathbf{\Gamma}^d\}$.
2. Next, sample from $\mathbb{P}(\tilde{\mathbf{M}}_{\leq 0} \mid C_{\leq 0})$. For $h > 0$, $\tilde{\mathbf{M}}_{t+h}$ can be obtained from $\tilde{\mathbf{M}}_t$ and $C_{\leq 0}$ by superimposing two Poisson point processes, conditionally independent given $C_{\leq 0}$:
 - (a) a point process of directed edges between lineages (the copying events), with rate $\frac{1}{2\eta_d}$ for $a \rightarrow a$, $b \rightarrow b$, $c \rightarrow c$ edges, and rate $\frac{1}{\eta_d}$ for $c \rightarrow a$, $c \rightarrow b$ edges.
 - (b) a point process of mutations hitting the lineages, at rate $\frac{\theta}{2}$ per locus per lineage.

To see that this superpositioning of point processes yields the correct distribution $\mathbb{P}(\tilde{\mathbf{M}}_{\leq 0})$, note that mutation and copying events do not affect the rates of recombination and recoalescence events, and that the 4 types of jump events that make up the Markov jump process $\tilde{\mathbf{M}}_{\leq 0}$ occur at the desired rates given by $\{\mathbf{\Lambda}^d\}$.

Now define C_t^* to be the *backwards in time* Markov chain with rates $\mathbf{\Gamma}^d$ in $(t_d, t_{d+1}]$ (whereas C_t has the same rates but going *forwards in time*). Let $\tilde{\mathbf{M}}_t^*$ be the stochastic process with conditional law $\mathbb{P}(\tilde{\mathbf{M}}_{\leq 0}^* \mid C_{\leq 0}^* = \mathcal{C}) = \mathbb{P}(\tilde{\mathbf{M}}_{\leq 0} \mid C_{\leq 0} = \mathcal{C})$. Thus $\tilde{\mathbf{M}}_{\leq 0}^*$ can be sampled in the same two steps as $\tilde{\mathbf{M}}_{\leq 0}$, except the first step (coalescences and recombinations) is backwards in time. We illustrate the conditional independence structure of $\tilde{\mathbf{M}}_t^*$ and C_t^* via a directed graphical model (Koller and Friedman, 2009) in Figure 11. (A graphical model is a graph whose vertices represent random variables, with the property that if all paths between V_1 and V_2 pass through W , then there is conditional independence $\mathbb{P}(V_1, V_2 \mid W) = \mathbb{P}(V_1 \mid W)\mathbb{P}(V_2 \mid W)$).




Figure 11: Probabilistic graphical model for the processes C_t^* and \tilde{M}_t^* , with $-\infty < s < t \leq 0$. Random variables are represented as vertices, and the edges encode conditional independence relationships. Specifically, if all paths between V_1 and V_2 pass through W , then there is conditional independence $\mathbb{P}(V_1, V_2 \mid W) = \mathbb{P}(V_1 \mid W)\mathbb{P}(V_2 \mid W)$.

We next show that for \mathbf{n} with $n^{(a)} = n^{(b)} = n - n^{(c)}$,

$$\mathbb{P}_t(\mathbf{n}) = \mathbb{P}(\tilde{\mathbf{M}}_t^* = \mathbf{n} \mid C_t^* = n^{(c)}). \quad (10)$$

We use a similar argument as in Durrett (2008, Theorem 1.30, p.47), tracing the genealogy of \mathbf{n} backwards in time (Figure 3b). Under $\mathbb{P}(\tilde{\mathbf{M}}^*)$, recombination events occur backwards in time at rate $\frac{\rho}{2}$ per c type lineage, as in the coalescent. Likewise, coalescence between an a and b type occurs at the usual rate $\frac{1}{\eta_d}$. Next, note that copying events between ancestral lineages induce coalescences within the ARG; these are encountered as a Poisson point process at rate $\frac{1}{\eta_d}$ per ancestral pair not of type (a, b) . Thus, the embedded ARG is distributed as the coalescent with recombination. Finally, conditioning on the full history of recombination, copying, and coalescence events, we can drop down mutations as a Poisson point process with rate $\frac{\theta}{2}$ per locus per lineage, and so the ARG with mutations follows the coalescent with recombination and mutation. Note that the alleles at the common ancestors of each locus follow the stationary distribution: the common ancestors are fixed under the conditioning (of recombination, copying, and coalescence events), and if \mathbf{v}_s is the conditional distribution of an ancestral allele at time $s \leq T_{\text{MRCA}}$, then $\mathbf{v}_s = \mathbf{v}_{s'} e^{(\mathbf{P} - \mathbf{I}) \frac{\theta}{2}(s-s')}$ for $s' < s$; sending $s' \rightarrow -\infty$ yields the stationary distribution.

Having established (10), we next observe

$$\begin{aligned} \mathbb{P}_{t_{d+1}}(\mathbf{n}) &= \mathbb{P}(\tilde{\mathbf{M}}_{t_{d+1}}^* = \mathbf{n} \mid C_{t_{d+1}}^* = n^{(c)}) \\ &= \sum_{\mathbf{m}} \mathbb{P}(\tilde{\mathbf{M}}_{t_d}^* = \mathbf{m} \mid C_{t_d}^* = m^{(c)}) \mathbb{P}(C_{t_d}^* = m^{(c)} \mid C_{t_{d+1}}^* = n^{(c)}) \\ &\quad \times \mathbb{P}(\tilde{\mathbf{M}}_{t_{d+1}}^* = \mathbf{n} \mid C_{t_{d+1}}^* = n^{(c)}, C_{t_d}^* = m^{(c)}, \tilde{\mathbf{M}}_{t_d}^* = \mathbf{m}) \\ &= \sum_{\mathbf{m}} \mathbb{P}_{t_d}(\mathbf{m}) \mathbb{P}(C_{t_d}^* = m^{(c)} \mid C_{t_{d+1}}^* = n^{(c)}) \\ &\quad \times \mathbb{P}(\tilde{\mathbf{M}}_{t_{d+1}}^* = \mathbf{n} \mid C_{t_{d+1}}^* = n^{(c)}, C_{t_d}^* = m^{(c)}, \tilde{\mathbf{M}}_{t_d}^* = \mathbf{m}). \end{aligned} \quad (11)$$

Note that in the second equality, we use the conditional independence of $\tilde{\mathbf{M}}_{t_d}^*$ and $C_{t_{d+1}}^*$ given $C_{t_d}^*$, which follows from the graphical model of Figure 11 by setting $s = t_d$ and $t = t_{d+1}$.

Next, note that $\mathbf{\Gamma}^d$ is the transition matrix of a simple random walk with bounded state space and no absorbing states, and thus is reversible. Thus, with γ^d the stationary distribution of $\mathbf{\Gamma}^d$,

$$\begin{aligned}\gamma_{n^{(c)}}^d \mathbb{P}(C_{t_d}^* = m^{(c)} \mid C_{t_{d+1}}^* = n^{(c)}) &= \gamma_{n^{(c)}}^d \left[e^{\mathbf{\Gamma}^d(t_{d+1}-t_d)} \right]_{n^{(c)}, m^{(c)}} \\ &= \gamma_{m^{(c)}}^d \left[e^{\mathbf{\Gamma}^d(t_{d+1}-t_d)} \right]_{m^{(c)}, n^{(c)}} \\ &= \gamma_{m^{(c)}}^d \mathbb{P}(C_{t_{d+1}} = n^{(c)} \mid C_{t_d} = m^{(c)}).\end{aligned}\tag{12}$$

Recall that we defined $\tilde{\gamma}_{\mathbf{n}}^d = \gamma_{n^{(c)}}^d$. So plugging (12) into (11) yields

$$\begin{aligned}\mathbb{P}_{t_{d+1}}(\mathbf{n}) &= \sum_{\mathbf{m}} \mathbb{P}_{t_d}(\mathbf{m}) \frac{\tilde{\gamma}_{\mathbf{m}}^d}{\tilde{\gamma}_{\mathbf{n}}^d} \mathbb{P}(C_{t_{d+1}} = n^{(c)} \mid C_{t_d} = m^{(c)}) \\ &\quad \times \mathbb{P}(\tilde{\mathbf{M}}_{t_{d+1}} = \mathbf{n} \mid C_{t_{d+1}} = n^{(c)}, C_{t_d} = m^{(c)}, \tilde{\mathbf{M}}_{t_d} = \mathbf{m}) \\ &= \sum_{\mathbf{m}} \mathbb{P}_{t_d}(\mathbf{m}) \frac{\tilde{\gamma}_{\mathbf{m}}^d}{\tilde{\gamma}_{\mathbf{n}}^d} \mathbb{P}(C_{t_{d+1}} = n^{(c)}, \tilde{\mathbf{M}}_{t_{d+1}} = \mathbf{n} \mid C_{t_d} = m^{(c)}, \tilde{\mathbf{M}}_{t_d} = \mathbf{m}) \\ &= \sum_{\mathbf{m}} \mathbb{P}_{t_d}(\mathbf{m}) \frac{\tilde{\gamma}_{\mathbf{m}}^d}{\tilde{\gamma}_{\mathbf{n}}^d} \left[e^{\tilde{\mathbf{\Lambda}}^d(t_{d+1}-t_d)} \right]_{\mathbf{m}, \mathbf{n}}.\end{aligned}$$

which proves half of the desired result, i.e. $\mathbf{p}^{d+1} = [(\mathbf{p}^d \odot \tilde{\gamma}^d) e^{\tilde{\mathbf{\Lambda}}^d(t_{d+1}-t_d)}] \div \tilde{\gamma}^d$, where $\mathbf{p}^d = [\mathbb{P}_{t_d}(\mathbf{n})]_{\mathbf{n}}'$. To show the other half, that $\mathbf{p}^{-D+1} = \tilde{\mathbf{\lambda}}^{-D} \div \tilde{\gamma}^{-D}$, where $\tilde{\mathbf{\lambda}}^{-D}$ is the stationary distribution of $\mathbf{\Lambda}^{-D}$, we simply note that for all $t \leq t_{-D+1}$,

$$\begin{aligned}\mathbb{P}_t(\mathbf{n}) \tilde{\gamma}_{\mathbf{n}}^{-D} &= \mathbb{P}(\tilde{\mathbf{M}}_t^* = \mathbf{n} \mid C_t^* = n^{(c)}) \gamma_{n^{(c)}}^{-D} \\ &= \mathbb{P}(\tilde{\mathbf{M}}_t = \mathbf{n} \mid C_t = n^{(c)}) \gamma_{n^{(c)}}^{-D} \\ &= \mathbb{P}(\tilde{\mathbf{M}}_t = \mathbf{n}) \\ &= \tilde{\lambda}_{\mathbf{n}}^{-D},\end{aligned}$$

where the second equality follows by reversibility of $\mathbf{\Gamma}^{-D}$, which implies $\mathbb{P}(C_{\leq t} \mid C_t) = \mathbb{P}(C_{\leq t}^* \mid C_t^*)$, and thus $\mathbb{P}(\tilde{\mathbf{M}}_{\leq t} \mid C_t) = \mathbb{P}(\tilde{\mathbf{M}}_{\leq t}^* \mid C_t^*)$.

D.2 Proof of Theorem 2

We first check that $\mathbb{P}(\mathbf{n}_{s_1} \mid \mathbf{n}_{s_2}, \mathbf{n}_{s_3}) = \mathbb{P}(\mathbf{n}_{s_1} \mid \mathbf{n}_{s_2})$, for $-\infty < s_1 < s_2 < s_3 \leq 0$, and so \mathbf{n}_t is a backwards in time Markov chain.

Recall that we generate $n_t^{(abc)}$ as a backwards in time Markov chain, then generate \mathbf{n}_t by dropping down mutations forward in time. The conditional independence structure of $\mathbf{n}_{s_1}, \mathbf{n}_{s_2}, \mathbf{n}_{s_3}$ is thus described by the directed graphical model (Koller and Friedman, 2009) in Figure 12.

Doing moralization and variable elimination (Koller and Friedman, 2009) on Figure 12 results in the undirected graphical model in Figure 13. The graphical model of Figure 13 then implies

$$\begin{aligned}\mathbb{P}(\mathbf{n}_{s_1} \mid \mathbf{n}_{s_2}, \mathbf{n}_{s_3}) &= \sum_{n_{s_2}^{(abc)}} \mathbb{P}(\mathbf{n}_{s_1} \mid \mathbf{n}_{s_2}, n_{s_2}^{(abc)}) \mathbb{P}(n_{s_2}^{(abc)} \mid \mathbf{n}_{s_2}, \mathbf{n}_{s_3}) \\ &= \mathbb{P}(\mathbf{n}_{s_1} \mid \mathbf{n}_{s_2}),\end{aligned}$$




Figure 12: Probabilistic graphical model for the coalescent with recombination and mutation, with $-\infty < s_1 < s_2 < s_3 \leq 0$.




Figure 13: Undirected graphical model, after moralization and variable elimination on Figure 12: we add edges to form cliques on the left and right sides of $n_{s_2}^{(abc)}$, \mathbf{n}_{s_2} , and then eliminate all the variables except the ones pictured here.

where the second equality follows because $n_{s_2}^{(abc)}$ is a deterministic function of \mathbf{n}_{s_2} . Thus, \mathbf{n}_t is a backwards in time Markov chain.

We next compute the backwards in time rates $q_{\mathbf{n}, \mathbf{m}}^{(t)}$ for the Markov chain \mathbf{n}_t at time t . Starting from the definition of $q_{\mathbf{n}, \mathbf{m}}^{(t)}$,

$$\begin{aligned}
q_{\mathbf{n}, \mathbf{m}}^{(t)} &= \left. \frac{d}{ds} \mathbb{P}(\mathbf{n}_{t-s} = \mathbf{m} \mid \mathbf{n}_t = \mathbf{n}) \right|_{s=0} \\
&= \left. \frac{d}{ds} \frac{\mathbb{P}(\mathbf{n}_{t-s} = \mathbf{m}, \mathbf{n}_t = \mathbf{n} \mid n_t^{(abc)} = n^{(abc)})}{\mathbb{P}(\mathbf{n}_t = \mathbf{n} \mid n_t^{(abc)} = n^{(abc)})} \right|_{s=0} \\
&= \frac{1}{\mathbb{P}_t(\mathbf{n})} \frac{d}{ds} \left[\mathbb{P}(n_{t-s}^{(abc)} = m^{(abc)} \mid n_t^{(abc)} = n^{(abc)}) \mathbb{P}(\mathbf{n}_t = \mathbf{n} \mid n_t^{(abc)} = n^{(abc)}, \mathbf{n}_{t-s} = \mathbf{m}) \mathbb{P}_{t-s}(\mathbf{m}) \right] \Big|_{s=0} \\
&= \frac{1}{\mathbb{P}_t(\mathbf{n})} \left[\mathbb{P}(n_t^{(abc)} = m^{(abc)} \mid n_t^{(abc)} = n^{(abc)}) \mathbb{P}(\mathbf{n}_t = \mathbf{n} \mid n_t^{(abc)} = n^{(abc)}, \mathbf{n}_t = \mathbf{m}) \frac{d}{ds} \mathbb{P}_{t-s}(\mathbf{m}) \Big|_{s=0} \right. \\
&\quad \left. + \phi_{\mathbf{n}, \mathbf{m}}^{(t)} \mathbb{P}_t(\mathbf{m}) \right] \\
&= \begin{cases} \phi_{\mathbf{n}, \mathbf{m}}^{(t)} \frac{\mathbb{P}_t(\mathbf{m})}{\mathbb{P}_t(\mathbf{n})}, & \text{if } \mathbf{m} \neq \mathbf{n}, \\ \phi_{\mathbf{n}, \mathbf{n}}^{(t)} - \frac{d}{dt} \log \mathbb{P}_t(\mathbf{n}), & \text{if } \mathbf{m} = \mathbf{n}, \end{cases}
\end{aligned}$$

where the penultimate equality follows from the product rule and the definition of $\phi^{(t)}$ in (5).

The specific entries of $\phi^{(t)}$ listed in Table 3 can be obtained by applying the product rule to (5), and noting that $\frac{d}{ds} \mathbb{P}(n_{t-s}^{(abc)} \mid n_t^{(abc)})|_{s=0}$ and $\frac{d}{ds} \mathbb{P}(\mathbf{n}_t \mid n_t^{(abc)}, \mathbf{n}_{t-s})|_{s=0}$ are, respectively, the backwards in time rates of $n_t^{(abc)}$ (as listed in Table 1), and the forward in time rates for dropping mutations on \mathbf{n}_t .

References

- 1000 Genomes Project Consortium. 2010. A map of human genome variation from population-scale sequencing. *Nature*, **467**, 1061–1073.
- Al-Mohy, A. H. and Higham, N. J. 2011. Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM Journal on Scientific Computing*, **33** (2), 488–511.
- Auton, A. and McVean, G. 2007. Recombination rate estimation in the presence of hotspots. *Genome Research*, **17**,(8) 1219–1227.
- Auton, A., Fladel-Alon, A., Pfeifer, S., Venn, O., Ségurel, L., Street, T., Leffler, E. M., Bowden, R., Aneas, I., Broxholme, J., et al. 2012. A fine-scale chimpanzee genetic map from population sequencing. *Science*, **336**,(6078) 193–198.
- Auton, A., Li, Y. R., Kidd, J., Oliveira, K., Nadel, J., Holloway, J. K., Hayward, J. J., Cohen, P. E., Greally, J. M., Wang, J., et al. 2013. Genetic recombination is targeted towards gene promoter regions in dogs. *PLoS Genetics*, **9**,(12) e1003984.
- Auton, A., Myers, S., and McVean, G. Identifying recombination hotspots using population genetic data. arXiv preprint: <http://arxiv.org/abs/1403.4264>, March 2014.
- Baudat, F., Buard, J., Grey, C., Fladel-Alon, A., Ober, C., Przeworski, M., Coop, G., and de Massy, B. 2010. PRDM9 is a major determinant of meiotic recombination hotspots in humans and mice. *Science*, **327**, 836–840.
- Berg, I. L., Neumann, R., Lam, K. G., Sarbajna, S., Odenthal-Hesse, L., May, C. A., and Jeffreys, A. J. 2010. PRDM9 variation strongly influences recombination hot-spot activity and meiotic instability in humans. *Nature Genetics*, **42**,(10) 859–863.
- Bhaskar, A. and Song, Y. S. 2012. Closed-form asymptotic sampling distributions under the coalescent with recombination for an arbitrary number of loci. *Advances in Applied Probability*, **44**, 391–407. (PMC3409093).
- Chan, A. H., Jenkins, P. A., and Song, Y. S. 2012. Genome-wide fine-scale recombination rate variation in *Drosophila melanogaster*. *PLoS Genetics*, **8**,(12) e1003090.
- Chen, G. K., Marjoram, P., and Wall, J. D. 2009. Fast and flexible simulation of DNA sequence data. *Genome Res.*, **19**, 136–142.
- Choudhary, M. and Singh, R. 1987. Historical effective size and the level of genetic diversity in *drosophila melanogaster* and *drosophila pseudoobscura*. *Biochemical genetics*, **25**,(1-2) 41–51.
- De Iorio, M. and Griffiths, R. C. 2004. Importance sampling on coalescent histories. I. *Adv. Appl. Prob.*, **36**, 417–433.
- Dialdestoro, K., Sibbesen, J. A., Maretty, L., Raghwani, J., Gall, A., Kellam, P., Pybus, O. G., Hein, J., and Jenkins, P. A. 2016. Coalescent inference using serially sampled, high-throughput sequencing data from intra-host hiv infection. *Genetics*. ISSN 0016-6731. doi: 10.1534/genetics.115.177931. URL <http://www.genetics.org/content/early/2016/02/03/genetics.115.177931>.
- Donnelly, P. and Kurtz, T. G. 1999. Genealogical processes for fleming-viot models with selection and recombination. *Annals of Applied Probability*, **9**,(4) 1091–1148.

- Durrett, R. *Probability Models for DNA Sequence Evolution*. Springer, New York, 2nd edition, 2008.
- Ethier, S. N. and Griffiths, R. C. 1990. On the two-locus sampling distribution. *J. Math. Biol.*, **29**, 131–159.
- Ethier, S. N. and Kurtz, T. G. 1993. Fleming-viot processes in population genetics. *SIAM Journal on Control and Optimization*, **31**,(2) 345–386.
- Fearnhead, P. 2003. Consistency of estimators of the population-scaled recombination rate. *Theoretical Population Biology*, **64**, 67–79.
- Fearnhead, P. and Donnelly, P. 2001. Estimating recombination rates from population genetic data. *Genetics*, **159**, 1299–1318.
- Fearnhead, P. and Smith, N. G. C. 2005. A novel method with improved power to detect recombination hotspots from polymorphism data reveals multiple hotspots in human genes. *Am. J. Hum. Genet.*, **77**, 781–794.
- Fearnhead, P., Harding, R. M., Schneider, J. A., Myers, S., and Donnelly, P. 2004. Application of coalescent methods to reveal fine-scale rate variation and recombination hotspots. *Genetics*, **167**, 2067–2081.
- Fearnhead, P. 2006. SequenceLDhot: detecting recombination hotspots. *Bioinformatics*, **22**, 3061–3066.
- Golding, G. B. 1984. The sampling distribution of linkage disequilibrium. *Genetics*, **108**, 257–274.
- Griffiths, R. C. and Marjoram, P. An ancestral recombination graph. In Donnelly, P. and Tavaré, S., editors, *Progress in population genetics and human evolution*, volume 87, pages 257–270. Springer-Verlag, Berlin, 1997.
- Griffiths, R. C., Jenkins, P. A., and Song, Y. S. 2008. Importance sampling and the two-locus model with subdivided population structure. *Advances in Applied Probability*, **40**, 473–500.
- Griffiths, R. 1991. The two-locus ancestral graph. *Selected Proceedings of the Sheffield Symposium on Applied Probability. IMS Lecture Notes–Monograph Series*, **18**, 100–117.
- Gutenkunst, R. N., Hernandez, R. D., Williamson, S. H., and Bustamante, C. D. 2009. Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genetics*, **5**,(10) e1000695.
- Hobolth, A., Uyenoyama, M., and Wiuf, C. 2008. Importance sampling for the infinite sites model. *Statistical Applications in Genetics and Molecular Biology*, **7**,(1) 32.
- Hudson, R. R. 1985. Sampling distribution of linkage disequilibrium under an infinite allele model without selection. *Genetics*, **109**, 611–631.
- Hudson, R. 2001. Two-locus sampling distributions and their application. *Genetics*, **159**,(4) 1805–1817.
- Jenkins, P. A. and Song, Y. S. 2009. Closed-form two-locus sampling distributions: accuracy and universality. *Genetics*, **183**, 1087–1103.

- Jenkins, P. A. and Song, Y. S. 2010. An asymptotic sampling formula for the coalescent with recombination. *Annals of Applied Probability*, **20**, 1005–1028. (PMC2910927).
- Jenkins, P. 2012. Stopping-time resampling and population genetic inference under coalescent models. *Statistical Applications in Genetics and Molecular Biology*, **11**,(1) 1–20. (PMC3800802).
- Jenkins, P. A. and Song, Y. S. 2012. Padé approximants and exact two-locus sampling distributions. *Annals of Applied Probability*, **22**, 576–607. (PMC3685441).
- Johnson, P. and Slatkin, M. 2009. Inference of microbial recombination rates from metagenomic data. *PLoS Genetics*, **5**,(10) e1000674.
- Johnston, H. R. and Cutler, D. J. 2012. Population demographic history can cause the appearance of recombination hotspots. *The American Journal of Human Genetics*, **90**,(5) 774–783.
- Kamm, J. A., Terhorst, J., and Song, Y. S. 2016. Efficient computation of the joint sample frequency spectra for multiple populations. *Journal of Computational and Graphical Statistics*, pages 1–37. doi: 10.1080/10618600.2016.1159212. URL <http://dx.doi.org/10.1080/10618600.2016.1159212>.
- Koller, D. and Friedman, N. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Koskela, J., Jenkins, P. A., and Spano, D. 2015. Computational inference beyond kingman’s coalescent. *Journal of Applied Probability*, **52**,(2) 519–537.
- Maruyama, T. Stochastic integrals and their application to population genetics. In Kimura, M., editor, *Molecular Evolution, Protein Polymorphism and their Neutral Theory*, pages 151–166. Springer-Verlag, Berlin, 1982.
- McVean, G., Awadalla, P., and Fearnhead, P. 2002. A coalescent-based method for detecting and estimating recombination from gene sequences. *Genetics*, **160**, 1231–1241.
- McVean, G. A. T. 2002. A genealogical interpretation of linkage disequilibrium. *Genetics*, **162**, 987–991.
- McVean, G., Myers, S., Hunt, S., Deloukas, P., Bentley, D., and Donnelly, P. 2004. The fine-scale structure of recombination rate variation in the human genome. *Science*, **304**,(5670) 581–584.
- Moran, P. 1 1958. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, **54**, 60–71. ISSN 1469-8064. doi: 10.1017/S0305004100033193.
- Myers, S., Bottolo, L., Freeman, C., McVean, G., and Donnelly, P. 2005. A fine-scale map of recombination rates and hotspots across the human genome. *Science*, **310**,(5746) 321–324.
- Myers, S., Bowden, R., Tumian, A., Bontrop, R., Freeman, C., MacFie, T., McVean, G., and Donnelly, P. 2010. Drive against hotspot motifs in primates implicates the PRDM9 gene in meiotic recombination. *Science*, **327**,(5967) 876–879.
- Myers, S., Freeman, C., Auton, A., Donnelly, P., and McVean, G. 2008. A common sequence motif associated with recombination hot spots and genome instability in humans. *Nature Genetics*, **40**,(9) 1124–1129.

- Ohta, T. and Kimura, M. 1969. Linkage disequilibrium due to random genetic drift. *Genet. Res. Camb.*, **13**, 47–55.
- Sheehan, S., Harris, K., and Song, Y. S. 2013. Estimating variable effective population sizes from multiple genomes: A sequentially Markov conditional sampling distribution approach. *Genetics*, **194**,(3) 647–662. (PMC3697970).
- Smith, N. G. C. and Fearnhead, P. 2005. A comparison of three estimators of the population-scaled recombination rate: Accuracy and robustness. *Genetics*, **171**,(4) 2051–2062. ISSN 0016-6731. doi: 10.1534/genetics.104.036293. URL <http://www.genetics.org/content/171/4/2051>.
- Song, Y. S. and Song, J. S. 2007. Analytic computation of the expectation of the linkage disequilibrium coefficient r^2 . *Theoretical Population Biology*, **71**, 49–60.
- Steinrücken, M., Kamm, J. A., and Song, Y. S. Inference of complex population histories using whole-genome sequences from multiple populations. bioRxiv preprint: <http://dx.doi.org/10.1101/026591>, September 2015.
- Stephens, M. and Donnelly, P. 2000. Inference in molecular population genetics. *J.R. Stat. Soc. Ser. B*, **62**, 605–655.
- Tajima, F. 1983. Evolutionary relationship of DNA sequences in finite populations. *Genetics*, **105**, (2) 437–460.
- The International HapMap Consortium. 2007. A second generation human haplotype map of over 3.1 million SNPs. *Nature*, **449**,(7164) 851–861.
- Wegmann, D., Kessner, D. E., Veeramah, K. R., Mathias, R. A., Nicolae, D. L., Yanek, L. R., Sun, Y. V., Torgerson, D. G., Rafaels, N., Mosley, T., Becker, L. C., Ruczinski, I., Beaty, T. H., Kardia, S. L. R., Meyers, D. A., Barnes, K. C., Becker, D. M., Freimer, N. B., and Novembre, J. 2011. Recombination rates in admixed individuals identified by ancestry-based inference. *Nat. Genet.*, **43**, 847–853.
- Weir, B. *Genetic data analysis II: Methods for discrete population genetic data*. Sinauer Associates, Sunderland, MA, 1996.
- Ye, M., Nielsen, S., Nicholson, M., Teh, Y., Jenkins, P., Colchester, F., Anderson, J., and Hein, J. Importance sampling under the coalescent with times and variable population. URL https://www.stats.ox.ac.uk/__data/assets/pdf_file/0007/9889/Coalescent_Sampling_Report.pdf. 2013.